Dissertation

submitted to the

Combined Faculties for the Natural Sciences and for Mathematics

of the Ruperto-Carola University of Heidelberg, Germany

for the degree of

Doctor of Natural Sciences

presented by:

Diplom-Physicist: Christian Thieke born in: Hann. Münden

Oral examination: 29th January, 2003

Multicriteria Optimization in Inverse Radiotherapy Planning

Referees: Prof. Dr. Thomas Bortfeld Prof. Dr. Josef Bille

Zusammenfassung

Multikriterielle Optimierung in der inversen Strahlentherapieplanung

Die Entwicklung der intensitätsmodulierten Strahlentherapie (IMRT) hat die Möglichkeiten zur Gestaltung der Dosisverteilung stark erweitert. Aber auch mit IMRT ist aufgrund physikalischer Begrenzungen eine gewisse Dosis in Risikoorganen unvermeidbar, so dass in der Bestrahlungsplanung oft ein Kompromiss zwischen einzelnen Plankriterien geschlossen werden muss. Diese Arbeit stellt ein neues Optimierungskonzept vor, das diesem Problemcharakter angepasst ist. Es basiert auf der Tatsache, dass es nicht einen einzigen optimalen Plan gibt, sondern viele sog. Pareto optimale Pläne, die in keinem Kriterium zu verbessern sind, ohne ein anderes zu verschlechtern. Es wird eine Datenbank mit allen Pareto optimalen Plänen erstellt und dem Planer die endgültige Entscheidung über den klinisch sinnvollsten Kompromiss überlassen. Wegen des erhöhten Rechenaufwandes wurde zunächst die Dosisberechnung des vorhandenen Planungssystems durch statistisches Ausdünnen der Berechnungsmatrizen um den Faktor drei beschleunigt. Zur Beschreibung eines Plans wird ein Vektor mit den 'equivalent uniform doses' (EUDs) der einzelnen Strukturen verwendet. Dazu wurden zwei Modelle der EUD analysiert, das sog. 'max & mean' und das 'generalized mean' Modell. Weiterhin wurde ein Verfahren entwickelt, das das 'generalized mean' Modell der EUD in die Optimierung eines Plans einbezieht. Zum Abschluss der Arbeit wurde für einen klinischen Beispielfall mit 1 Zielvolumen und 2 Risikoorganen die Datenbank aller Pareto optimalen Lösungen generiert und visualisiert. Der Planer kann damit auf einen Blick die Abhängigkeiten der einzelnen Planstrukturen erkennen und sich für den klinisch besten Plan entscheiden.

Abstract

Multicriteria Optimization in Inverse Radiotherapy Planning

The development of intensity-modulated radiotherapy (IMRT) has greatly enhanced the possibilities of shaping the dose distribution. However, because of physical limitations, even with IMRT a certain dose to organs at risk is unavoidable, so in treatment planning often a compromise between different plan criteria has to be found. This thesis presents a new optimization concept that is dedicated to this type of problem. It is based on the fact that there is no single optimal plan, but many so-called Pareto optimal plans, which cannot be improved in one criterion without worsening another. A database with all Pareto optimal plans is generated and the decision which compromise is clinically the best is left to the planner. Because of the higher calculation complexity, the dose calculation of the existing planning system was accelerated by a factor of three through statistical sampling of the calculation matrices. To describe a plan, a vector containing the 'equivalent uniform doses' (EUDs) of each structure is used. Therefore two EUD models were analyzed, the 'max & mean' and the 'generalized mean' model. An algorithm was developed to include the 'generalized mean' model in the plan optimization. At the end of the thesis, the database containing all Pareto optimal solutions of a clinical test case with 1 target volume and 2 organs at risk was generated and visualized. The planner can see at first sight the dependencies of the different plan elements and can decide which plan is clinically the best.

Contents

1	Intr	oduction	1			
2	The multicriteria approach in intensity-modulated radiotherapy -					
	A new optimization paradigm					
	2.1	Current optimization strategies	6			
	2.2	The new optimization paradigm	9			
		2.2.1 Database generation	10			
		2.2.2 Interactive planning session	12			
3	Acceleration of dose calculation through importance sampling					
	3.1	Introduction	16			
	3.2	Methods	18			
	3.3	Results and Discussion	23			
4	The equivalent uniform dose: Models and parameters					
	4.1	Introduction	31			
	4.2	The generalized mean model	33			
	4.3	The max & mean model \ldots	34			
	4.4	Results	37			
	4.5	Discussion	39			
5	The equivalent uniform dose as an optimization constraint					
	5.1	Introduction	44			
	5.2	Material and Methods	45			
	5.3	Results	52			
	5.4	Discussion	57			

6	Practical application of multicriteria optimization in inverse ra-							
	diotherapy planning 6							
	6.1 Clinical test case with 1 target volume and 2 organs at risk 6							
		6.1.1	Introduction	63				
		6.1.2	Material and Methods	64				
		6.1.3	Results and Discussion	65				
	6.2 Techniques for more complex cases							
		6.2.1	Adaptive constraints	69				
		6.2.2	Plan interpolation	70				
7	Summary and Outlook 73							

Chapter 1

Introduction

Radiotherapy is, after surgery, the second most important treatment option against cancer. Every second cancer patient is treated with radiation, as the only form of treatment or in combination with surgery and chemotherapy. The goal of radiotherapy is to achieve tumor control without causing complications. Therefore a high dose of radiation must be delivered to the tumor and, at the same time, the dose to the surrounding normal tissue must be as low as possible.

A main research topic in radiotherapy is focussed on realizing a high degree of dose conformality, which is the restriction of the high dose area to the tumor target volume. During the last decade, a major improvement in this field was achieved by the development of intensity-modulated radiotherapy (IMRT). As the name implies, in IMRT the intensities of the incident beams can be varied and are no longer uniform over the whole aperture. Using this technique, complex shaped dose distributions can be realized. This allows better sparing of organs at risks and higher doses to the target volume, especially in cases where the target volume is concave and directly adjacent to organs at risk. Nevertheless, because of physical limitations, even IMRT cannot deliver the ideal dose distribution, which is 100% in the target volume and 0% everywhere else. While IMRT does allow to achieve somewhat steeper dose gradient between normal tissue and the target volume, especially for irregularly shaped targets, the steepness of the gradients is still physically limited; the 20% to 80% penumbra is at least about 6 mm wide. As a consequence of this, if the target is directly abutting to a critical structure, the minimum target dose equals the maximal dose in the critical structure. In a

more general sense, the goals of delivering a high dose to the target and a small dose to the critical structures contradict each other. A realistic objective can only be to find a suitable compromise between target coverage and normal tissue sparing. It is part of the treatment planning to find the compromise which is clinically the best for the individual patient.

In conventional radiotherapy forward planning, the planner usually defines the treatment setup (e.g. number and directions of beams, wedges, apertures), lets the computer calculate the resulting dose and changes the setup until the plan is acceptable. This is no longer possible in IMRT because of the many degrees of freedom - the intensities of up to several hundred beam elements have to be adjusted independently. Therefore the inverse planning scheme is used: The planner prescribes what dose inside the target volume is desired and what doses in critical structures can be tolerated, and the computer tries to find the treatment plan which is optimal with respect to the prescriptions. Most current inverse planning systems accomplish this task through an iterative optimization process that grades the plans by a one-dimensional objective function, so the optimal plan is the one with the best objective function value. It is quite obvious that characterizing something as complex as a treatment plan by a single number is extremely difficult. Often the resulting plan is clinically not satisfactory, and repeated optimization runs with different constraint settings may become necessary. This trial and error process is time consuming and holds the risk of clinical sub-optimal treatment plans.

This thesis is focussed on a different optimization strategy, in which IMRT optimization is tackled using a so-called multicriteria approach. As said, there are several criteria in a treatment plan which cannot be optimal all at the same time. This makes the plan optimization problem inherently multicriterial: There is no single best solution - instead, there are many best compromises, also called Pareto optimal or efficient. A solution is Pareto optimal if it cannot be improved in one criterion without worsening at least one of the others. As a measure for the dose inside each structure, the equivalent uniform dose (EUD) is used. It is defined as the homogenous dose with the same clinical effect as the given inhomogenous dose distribution. This is potentially of higher clinical relevance than other dose distribution parameters like the minimum and maximum physical energy dose.

In the multicriteria concept, a treatment plan is no longer represented by an abstract scalar objective function value, but rather as a vector containing the EUD values of each structure.

The goal of multicriteria optimization in treatment planning is to generate a database containing all Pareto optimal solutions that might be of clinical interest, and leave the decision which solution is eventually the best for the patient to the treatment planner. The dose distributions in the relevant critical structures and in the target volume are controlled separately and interactively. An intuitive graphical interface assists the planner in browsing the database.

The thesis is structured as follows. Chapter 2 discusses the current state of the art in inverse planning and the concept of multicriteria optimization in more detail. Chapter 3 addresses the acceleration of dose calculation inside the optimization loop. Chapter 4 introduces the equivalent uniform dose, compares two different EUD models and fits the models to clinical data. Chapter 5 describes how the EUD is implemented as an optimization constraint. In chapter 6 the multicriteria concept is used to optimize a simple clinical case. All algorithms presented so far are utilized, and some other techniques that will be necessary for more complex cases are introduced as well. Chapter 7 summarizes the presented material. It also includes an outlook for the future work on this project.

Chapter 2

The multicriteria approach in intensity-modulated radiotherapy - A new optimization paradigm

Intensity-modulated radiotherapy (IMRT) has a much greater potential to shape dose distributions than conventional radiotherapy with uniform beams (IMRT Collaborative Working Group 2001). This capability has been used to tailor the dose distribution to the tumor target volume in conformal radiotherapy. In general IMRT allows one to achieve a better dose conformality, especially for irregularly shaped concave target volumes. The improved physical characteristics of IMRT can lead to improved clinical results, as was suggested in recent clinical studies (Zelefsky et al. 2000). Although IMRT is already in clinical use at several hospitals in Europe and many in the USA, there is a lot of potential and need for further improvements. The first hurdle is the initial implementation and the commissioning of the IMRT system in a clinical environment. Vendors advertise turn-key IMRT solutions but the reality may look different. Often the problem is to link the different elements of the IMRT planning and delivery chain. Current developments aim at more streamlined and integrated solutions. Another important point is that - because IMRT is more complex - it requires an elaborate, if possible patient-specific verification. More efficient tools are needed to make this

process less labor-intensive and time-consuming. Most vendors now offer IMRTready multileaf collimators (MLC) that are able to deliver IMRT automatically, without user intervention. The delivery time varies between the vendors but is generally clinically acceptable. Other developments aim at improving the planning of IMRT. This involves at the far end the translation of calculated intensity maps into sequences of MLC positions. On the other hand there is the calculation of the intensity maps itself. Some recent approaches also try to optimize the sequence of MLC shapes directly, without going through the intermediate step of the intensity maps. In most systems IMRT planning is considered as an optimization problem. The goal is to find the parameters–intensity maps, sometimes also beam orientations (e.g. Pugachev 2001), energy, etc.–that yield the best possible treatment plan under consideration of various clinical, technical, and physical constraints. A huge number of current research activities is related to this problem.

This thesis is focussed on the optimization problem as well. In this chapter, the current status of existing systems and a new multicriteria approach (cf. Bortfeld et al. 2002) is presented.

2.1 Current optimization strategies

In current IMRT planning systems, following techniques exist to control the optimization result:

1. Weight factors

One way is to combine objectives or costlets for the different critical structures (F_R) and the target volume (F_T) using weight factors w (also called penalties or importance factors). Mathematically, an objective function is defined that is of the form $F = w_T F_T + w_{R1} F_{R1} + w_{R2} F_{R2} + \ldots + w_K F_K$. By using a large value for w_T , more emphasis is placed on the target dose, and vice versa. Figure 2.1 shows dose-volume histograms of optimized IMRT plans for the treatment of a head and neck tumor. Two extreme values of the weight factor were used. A problem with this approach is that the weight factors have no clinical meaning. Suitable weight factors have to be determined by trial and error, which may be quite time consuming.



Figure 2.1: IMRT planning for a head and neck tumor using optimization with weight factors. Plan 1 is the optimization result with weight factors 1000 for target and 10 for spinal cord and brainstem, whereas plan 2 is the result with weightings 10 for target and 1000 for spinal cord and brainstem. As can be seen from this figure, the artificial weighting factors can radically change the optimization result. The equivalent uniform dose (EUD, see chapter 4) for the target volume changes from 57.8 Gy for plan 1 to 3.2 Gy for plan 2.



Figure 2.2: Same case as in Figure 2.1. Here constrained optimization was used. The dose constraints for the spinal cord and the brainstem were set to 37 Gy. According to these constraints, plan 1 is mathematically equivalent to Plan 2. However, plan 2 is clearly better in terms of sparing the brainstem.

2. Constraints

Another approach, which is also used in commercial planning systems, uses constraints on the maximum dose in the critical structures or the minimum dose in the target. Here the potential problem is that there is no reward for reducing the dose in the critical structures below the tolerance. Figure 2.2 shows an example with two critical structures, one of which is dose limiting for the target volume, but the other one is not. In fact, it is easily possible to reduce the dose in the brainstem without compromising the dose in any of the other structures. Hence, although both plans in figure are clinically acceptable, plan 2 is clearly better. A problem with both of the above approaches (weighting and constraining) is that in general the sensitivity of the optimization result to changes is unknown. That means, it is unclear how dependent the dose in one structure is on the constraints in other structures, e. g. the planner does not know how much dose in a critical structure could be saved if he is willing to accept a certain target dose reduction.

After the first optimization run, the system will come up with a result that is mathematically optimal with respect to the given objective function and constraints. However, this does not mean that the plan is also clinically acceptable. The reason is that the optimization criteria used in the optimization may not be clinically relevant. If a physician looks at the plan he or she usually wants to change the plan here and there. Now the treatment planner has the task to perform the desired modification with the available "knobs", that is with dose constraints, weight factors, and so on. Often these knobs are not well suited to do the modification. Several runs are generally required to come up with a suitable solution. This trial and error can take a lot of time. This is why optimized "inverse" planning is not a push-button solution.

2.2 The new optimization paradigm

To solve the above-mentioned drawbacks and problems in IMRT planning, a new optimization concept was developed in collaboration between the German Cancer Research Center (DKFZ) in Heidelberg/Germany, the Fraunhofer Institute for Industrial Mathematics (ITWM) in Kaiserslautern/Germany, and the Massachusetts General Hospital (MGH) in Boston/USA. Instead of defining a scalar objective function as a measure of the quality of the treatment plan, the new approach is inherently multicriterial. The multicriteria optimization concept was already used in planning of brachytherapy (Lahanas et al. 1999, Milickovic et al. 2002), but up to now no solution existed for planning of external radiotherapy. The dose distribution in each structure (critical structures and target) is characterized by its own, separate parameter. There are several options of what this parameter can be. One option would be the sum of quadratic differences to a specified dose constraint, which is a common criterion in existing planning systems. In the following, a probably clinically more relevant parameter will be used, the equivalent uniform dose (EUD). The EUD is defined as the uniform dose that leads to the same clinical effect as the actual non-uniform dose in that organ. Using EUD, volume effects (most noticable in parallel organized organs like the lung) are automatically considered in the optimization process. The whole plan is characterized by a vector containing the EUDs of all structures, $\vec{F} = (F_T, F_{R1}, F_{R2}, \ldots, F_K)$. What we seek is an optimal compromise between the target and the critical structures. Now, there is an infinite number of feasible combinations of those parameter sets. To make the search tractable, we look at so-called efficient (Pareto optimal) solutions only. These are defined as solutions in which an improvement in one organ will always lead to a worse result in at least one of the other organs.

The advantages of this concept are threefold:

- 1. Artificial weight factors, which have no clinical meaning, are avoided. The whole concept is based on dose-like values, which are amenable to a clinical interpretation.
- 2. Unnecessarily high doses in some of the critical structures, which can occur in constrained optimization (see above), are avoided by definition of the efficient (Pareto optimal) solution.
- 3. Plan tuning can be done interactively using "knobs" that have a clinical meaning. It is easy to do a sensitivity analysis and determine the dependency of, say, the target EUD on any of the critical structure EUDs.

The optimization procedure is split into two parts: First, the system automatically generates a database of Pareto optimal plans. Once started, this step does not require user interaction and can be performed over night. The second part is the interactive planning session in which the planner explores the database and finally selects the plan that will be used for treatment.

2.2.1 Database generation

Goal of this stage is generating a database containing a *representative system* of all Pareto optimal plans that are of clinical interest. A graphical illustration

of the problem is shown in Figure 2.3. It is not practible to store *all* plans that are Pareto optimal: E.g., all intensities set to zero will lead to zero dose in every voxel, which is a Pareto optimal plan because due to scatter any dose improvement in the target will worsen the dose in critical structures. But a plan with no dose to the target is clearly not of clinical interest. Another extreme would be a plan with intolerable high doses to vital organs at risk like the brainstem. So, only a subset of all Pareto optimal plans is of clinical interest. Even with restriction to this subset, there is an infinite number of Pareto optimal plans left because the dose distributions continuously change with the intensities and the difference between two Pareto optimal plans can become infinitesimal small. If the differences or dose volume histograms, it would not make sense to store both of these plans in the database. Instead, all plans in the database should be significantly different (resolution criterion) and at the same time nearly equidistantly cover the whole area of clinical interest (homogeneity criterion).

In practice, the procedure goes as follows: Initially the treatment planner roughly specifies the desired dose in the target volume and the dose he is willing to tolerate inside the organs at risk. This will lead to a first treatment plan, which is stored as the starting solution. Subsequently the planning system will produce all other plans and build the database according to the criteria given above. Every plan will be normalized to the desired dose in the target volume, and for each plan following information is stored in the database:

- F-vector containing the EUD values of all structures
- treatment setup (e.g. number and directions of beams, intensity maps)
- data for dose visualisation (e.g. isodose vectors, dose volume histogram, 3D dose cube)
- relationship to the neighboring plans

The plans are calculated and the database is filled over night. One can see that in this concept the definition of the initial constraints is not critical. It is only used to generate a starting solution that lies somewhere in the area of clinical interest.



Area of clinical interest

Figure 2.3: Only a subset of all Pareto optimal solutions is of clinical interest. It is part of the multicriteria optimization concept to generate a database with representative solutions covering the area of clinical interest.

This thesis is mainly focussed on the database generation part of the multicriteria optimization concept. As development basis and calculation platform, an experimental inverse planning system (Nill 2001) based on KonRad (Preiser et al. 1997) was used. The algorithms which are presented in the next chapters were all integrated into and tested with this system. This way there was no need for re-programming basic functionalities like CT data import, beam configuration setup and visualization capabilities.

2.2.2 Interactive planning session

After the database has been created, the treatment planner can browse through it interactively. Figure 2.4 shows the graphical user interface of a search tool that was developed by mathematicians from the Fraunhofer Institute for Industrial Mathematics in Kaiserslautern. The screen is divided into four main windows: The information window shows basic information like patient name and irradiation device. The current plan is shown in the isodose and dose volume histogram windows. The central part of the program is the navigation window. Here, every structure of the plan is represented with its own axis used for the EUD values. In the example of Fig. 2.4 there are one target and three organs at risk. All EUD values that can be reached through the database are defining the planning horizon, shown in the navigation window as the area with gray background. The planner sees right from the start which EUD value in each organ can reached. The current plan is highlighted in the navigation window by a polyhedron (purple line fig. 2.4) defined by the EUD value of each structure. If one or more aspects of the plan are not desirable (e.g., the EUD in one critical structure is too high), the planner can immediately go to another plan which satisfy that specific criterion. This is done by clicking and dragging the EUD marker of the particular organ in the desired direction. Instantaneously, the system finds the corresponding plan in its database and updates the information for the other structures and in the isodose and dose volume histogram windows. This gives the planner immediate feedback about the sensitivity of the problem. To give more control over the plan selection, it is possible to "lock" an organ by clicking on the box at the end of the respective EUD bar. If an organ is "locked", all treatment plans with a worse EUD than the actual one are excluded from the further database exploration. The reduced planning horizon is visualized in the navigation window by a different shade of the gray background (see also fig. 2.4, where the organ at risk "h3" is locked).

As mentioned in the database generation section, all plans in the database have the same target EUD value. This is the reason why there is no "lock" box for the target bar. But nevertheless the target EUD can be changed. This is done in a different way than for organs at risk: By dragging the target EUD arrow, all plans in the database are re-scaled to the desired target EUD. All other plan components are scaled by the same factor, so the range of EUD values that can be reached in the organs at risk will change as well. The advantages of this concept are that there is no restriction for the target dose, and that for every target dose the complete database of Pareto optimal plans is accessible.

It is not necessary to look at each and every plan stored in the database. Instead, the interface is designed to guide the planner to the best plan for the patient in as few steps as possible.



Figure 2.4: Prototype version of the interactive search tool developed at the ITWM in Kaiserslautern. In the navigation window, every structure is represented by a bar indicating the range of feasible EUD values. The dose distribution of the actual treatment plan is also visualized as isodoses and as dose volume histogram.

Chapter 3

Acceleration of dose calculation through importance sampling

In inverse planning for intensity-modulated radiotherapy, the dose calculation is a crucial element limiting both the maximum achievable plan quality and the speed of the optimization process. One way to integrate accurate dose calculation algorithms into inverse planning (implemented in the inverse planning program used for this thesis) is to precalculate the dose contribution of each beam element to each voxel for unit fluence. These precalculated values are stored in a big dose calculation matrix. Then the dose calculation during the iterative optimization process consists merely of matrix look-up and multiplication with the actual fluence values. However, because the dose calculation matrix can become very large, this ansatz requires a lot of computer memory and is still very time-consuming, making it not practical for clinical routine without further modifications.

The time problem becomes even more serious when in a multicriteria setting many plans have to be created. As a prerequisite for the multicriteria optimization project, it was therefore indispensable to accelerate the dose calulation algorithm of the inverse planning system.

This chapter presents a new method to significantly reduce the number of entries in the dose calculation matrix. The method utilizes the fact that a photon pencil beam has a rapid radial dose fall-off, and has very small dose values for the most part. In this low-dose part of the pencil beam, the dose contribution to a voxel is only integrated into the dose calculation matrix with a certain probability. Normalization with the reciprocal of this probability preserves the total energy, even though many matrix elements are omitted. Three probability distributions are tested to find the most accurate one for a given memory size. The sampling method is compared with the use of a fully filled matrix and with the well-known method of just cutting off the pencil beam at a certain lateral distance. A clinical example of a head and neck case is presented. It turns out that a sampled dose calculation matrix with only 1/3 of the entries of the fully filled matrix does not sacrifice the quality of the resulting plans, whereby the cut-off method results in a sub-optimal treatment plan.

3.1 Introduction

Various inverse treatment planning programs are now commercially available. Common to all is that they use simplified dose calculation algorithms in the optimization of the treatment plans. Therefore, in most cases an independent, accurate dose calculation is recommended as the final step in inverse planning in order to verify the dose distribution. Although this is not a big burden, it makes the inverse planning process less efficient than it could be. More importantly, in geometrically or physically complex cases, e.g. with severe tissue inhomogeneities, the dose distribution calculated with inverse planning can deviate significantly from the more accurately calculated distribution, and could therefore be suboptimal.

An obvious approach to overcoming this problem is the use of more accurate dose calculation algorithms within the inverse planning program. However, because inverse planning is iterative and requires the dose to be calculated in the order of 100 or more times, this is in conflict with the requirement that inverse planning should be reasonably fast, and ideally interactive. In a recent publication it was reported that using a hybrid approach combining one simplified and one more accurate dose calculation, the number of accurate calculations could be significantly reduced down to 5-10 (Siebers et al. 2001). Nevertheless, this still can be quite time consuming.

The inverse planning system that was used as calculation platform for this thesis uses another method. It calculates the dose distribution only once for each pencil beam before the optimization, store it in a large matrix, and access this pre-calculated dose in each iteration step of the inverse planning procedure. The system knows the dose in each voxel, and it knows the contributions from each bixel to this dose. In addition to the accurate dose calculation, this has the advantage that the lateral scatter can be taken into account also in the projection step of the optimization where the bixel intensities are updated.

Fortunately, the increased RAM storage capacity of modern computers has almost made it possible to implement this approach practically. Unfortunately, we are not quite there yet. Consider a realistic case in which there are 500 beam elements (bixels) in each of 9 treatment fields, and the number of volume elements (voxels) for which the dose is to be optimized is 200,000. If each dose value is stored as a 2 byte integer value in combination with a 4 byte index to the voxel, the required RAM storage space is in the order of 4-5 Gbyte, which is still too much for most modern computers. Moreover, the calculation of all these matrix entries (even it is done only once) and their subsequent access in the optimization loop would be time consuming.

A simple way to reduce the required storage space and increase the calculation speed is to calculate and store dose values only for the critical structures and the target volume. This alone leads to a sevenfold speed increase, depending on the specific case. For a fast calculation of specific treatment plan parameters (e.g., the DVH or TCP/NTCP-values) it was proposed to calculate only a small subset of the dose points inside a region of interest (Niemierko and Goitein 1990, Lu and Chin 1993). A speedup of the complete dose calculation is possible with the help of certain approximations. The crude method is to cut the pencil beams at a certain radial distance, say 15 mm, i.e., all elements beyond 15 mm radius are ignored (Cho and Phillips 2001). However, this approach leads to obvious errors in the scatter dose distribution and in the output factor.

The goal of this chapter (see also Thieke et al. 2002a) is to develop compact storage and access techniques for pencil beams, particularly but not solely in inverse treatment planning. The principal idea is to sample the pencil beam dose distribution according to the pencil beam value. The central part of the pencil beam is fully sampled, while outer regions are sampled with a certain probability distribution, i.e., not every element is stored. A weighting is applied to conserve the imparted energy, even when a certain percentage of pencil beam elements is omitted. Because in general many pencil beams contribute to each voxel, the result is a correctly predicted dose distribution with a small statistical fluctuation. The approach is similar to Monte Carlo integration techniques.

3.2 Methods

Because dose is defined as the absorbed energy, the dose to a volume element (voxel) in the patient can be written as the linear superposition of dose contributions from each beam. This concept can be extended to IMRT, where the contribution from each beam element (bixel) has to be considered. If the IMRT is delivered with a multileaf collimator (MLC), non-linearities can be introduced by certain leaf designs, such as the tongue and groove effect (Mohan 1995). However, these effects are minor and are disregarded here. We therefore write the dose d_i at voxel i in the form

$$d_i = \sum_j D_{ij} w_j, \tag{3.1}$$

where w_j is the weight (proportional to fluence and intensity) of the *j*-th bixel (macro pencil beam). The dose calculation matrix D is a pre-calculated matrix where element D_{ij} is the fractional dose contribution from a normalized pencil beam *j* to the voxel *i*. Typical resolutions are $(1 \text{ cm})^2$ per bixel and $(2.5 \text{ mm})^3$ per dose voxel, resulting in an order of magnitude of $10^9 D_{ij}$ elements for a typical clinical case. If higher resolutions are required, the number of D_{ij} elements increases correspondingly. As mentioned in the introduction, a computer cannot handle this amount of data easily, even if the latest generation is used.

Let us consider a particular macro pencil beam with index j. We should recall that a pencil beam has a rapid lateral dose fall-off, i.e., the value of D_{ij} strongly depends on the radial distance r between the central ray of the pencil beam jand the voxel i. In fig. 3.1 the lateral dose distribution of a 15 MV $(1 \text{ cm})^2$ pencil beam in water is shown at a depth of 10 cm (solid line).



Figure 3.1: Dose distribution of a $(1 \text{ cm})^2$ 15 MV photon pencil beam in water at a depth z=10 cm.

The dose was calculated by a clinically approved pencil beam algorithm developed at the German Cancer Research Center by Bortfeld et al. (1993). At a radial distance r = 25 mm the dose has already dropped to less than 1% of the maximum value at the central ray. Therefore a simple approach for reducing the size of the matrix D is to omit all D_{ij} elements above a certain radial distance or below a certain dose value. However, the error introduced into the dose calculation by using such a reduced matrix is quite noticeable. This can be seen from the dashed curve in fig. 3.1, which represents the integral of the pencil beam up to the radial distance r and which can be considered as the output factor. If the lateral cut off distance is set to 25 mm, the calculated output factor does not increase for fields with radii larger than 25 mm. This can lead to errors of 16% for large fields. To keep the error of the dose calculation below 4%, the pencil beam must be taken into account up to a lateral distance of 82 mm. Wherever the term "fully filled D_{ij} matrix" is used in the following, it refers to a dose calculation matrix where all elements up to a radial distance of 82 mm from the voxel i to the central ray of the bixel j are accounted for.

In this chapter a new approach is proposed for omitting a large percentage of the D_{ij} elements without causing the problems associated with the cut-off approach. The idea is that, starting from a radial distance r_0 , we consider D_{ij} elements only with a certain probability P_{ij} . With a probability of $1 - P_{ij}$ the element D_{ij} will be disregarded. Elements within a circle of radius r_0 are considered with a probability of $P_{ij} = 1$. To preserve the total energy, the element D_{ij} is multiplied by P_{ij}^{-1} , i.e., instead of D_{ij} we use

$$\widetilde{D}_{ij} = [D_{ij}]_{P_{ij}} \cdot P_{ij}^{-1}.$$
(3.2)

The subscript P_{ij} indicates that the element D_{ij} is integrated into D_{ij} with the probability P_{ij} . This approach ensures that even outer regions of the pencil beams are represented, and the energy delivered by the beam is completely enclosed in the D_{ij} matrix. The tradeoff for reducing the number of matrix elements this way is random noise in the dose distribution, but it can be expected that the noise is not very high due to two reasons: First, the sampling is only done in low-dose regions, and second, usually for every voxel many pencil beams contribute dose, so many noise patterns are superimposed and the resulting noise level is lower than for one single pencil beam. For this to work effectively the random sampling has to be performed individually for each pencil beam; a fixed sampling pattern would introduce artefacts in the resulting dose distribution.

Both the number of the resulting \widetilde{D}_{ij} elements and the noise level depend on the radius r_0 and the probability distribution used for the sampling outside r_0 . Three types of probability distributions have been tested: constant sampling probability, linearly decreasing probability beginning with P = 1 at $r = r_0$, and a sampling probability that is proportional to the original pencil beam dose value at the particular voxel (also beginning with P = 1 at $r = r_0$). The latter leads to a constant value of \widetilde{D}_{ij} at radii larger than r_0 . All three modes of sampling are graphically compared in fig. 3.2.



Figure 3.2: Sampling of the $(1 \text{ cm})^2$ pencil beam with (a) constant, (b) linear decreasing, and (c) dose proportional sampling probability. The voxel resolution was set to 2 mm, the sampling starts at r = 25 mm. The probability distribution is shown as a dashed line. For each probability distribution, an exemplary set of nine sampling values is shown as dashed bars.

The correlation between the number of \widetilde{D}_{ij} elements and the accuracy of dose calculation was investigated for all three sampling modes, as well as for the cutoff technique. Accuracy means the degree of agreement between the calculations with the reduced matrices and those with the full dose calculation matrix. As stated above, nonlinear effects of the beam delivery device like the tongue and groove effect are disregarded in the D_{ij} approach in general. To find the optimal sampling parameters for clinical use, treatment plans (including beam profiles) were compared resulting from inverse planning with fully filled and reduced matrices. As a reference standard for calculating the accuracy of each method the approved pencil beam algorithm mentioned above was used. The existing inverse planning system (Nill 2001) was enhanced by the option of using the pencil beam sampling technique. In both cases—full and reduced matrices—the optimized fluences were transferred to the independent, approved dose algorithm (Bortfeld et al. 1993) and the dose was recalculated. In the case of the cutoff technique, the dose (and the intensity) was scaled to the correct mean target dose because, for reasons mentioned above, this techniques did not yield the correct absolute dose level. Several clinical cases (a head and neck tumor shown in fig. 3.3, a prostate cancer case and a skull base tumor) have been tested, and various settings for r_0 (20, 25 and 30 mm) and the cut-off value (35, 40, 45, and 50 mm) were used for the optimization to find out the maximum reduction of the dose calculation matrix not affecting the quality of the resulting test plans.



Figure 3.3: 3D model and exemplary CT slices of the tested clinical head and neck case.

3.3 Results and Discussion

In fig. 3.4a the dose profile of a 15 MV $(10 \text{ cm})^2$ photon beam in water at a depth of z=10 cm is shown, calculated by different matrices: first by a full D_{ij} matrix, and further by reduced \widetilde{D}_{ij} matrices with approx. 1/3 of the original size ($r_0=25 \text{ mm}$, sampling with dose-dependent and constant probability). Using linear falling sampling probability for the matrix reduction, the quality of the

results falls between the constant and the dose dependent sampling. As can be seen from the figure, dose dependent sampling leads to the lowest noise level for a given number of \widetilde{D}_{ij} elements. This fact was validated for different field sizes and \widetilde{D}_{ij} sizes, not shown here. These results can be explained by the different sampling densities and error distributions along the lateral axis (see also fig. 3.2): for dose dependent sampling, all \widetilde{D}_{ij} values outside r_0 are equal to the dose at $r = r_0$; the maximum error occurs at large radial distances with very small dose contributions. At the inner part of the sampling area, beginning at r_0 , the error is very small. Using a constant sampling probability, the error close to r_0 is higher; the denser sampling in the outer regions (in relation to the dose dependent sampling) cannot compensate for this. Sampling with strictly linear decreasing probability suffers from the fact that, for practical settings of r_0 and a fixed matrix reduction factor, the sampling probability drops to 0 before a radial distance of 82 mm has been reached. In the following only the dose dependent sampling will be used.



Figure 3.4: Dose profile of a $(10 \text{ cm})^2$ 15 MV photon beam in water at a depth z = 10 cm, calculated (a) by a full and by sampled matrices (constant and dose dependent sampling probability) and (b) by a full, a dose dependent sampled and a cut-off matrix.

Fig. 3.4b again shows the dose profiles for the full and the dose-dependent sampled matrices and compares them with the cutoff technique. The cutoff value was set to r = 45 mm, so the cutoff matrix equals the size of the sampled matrix. The profile resulting from the cutoff matrix is smooth, but consistently underestimates the dose. Because it equals 0 for all points further than 45 mm from the geometrical field border, the underestimation of the scatter cannot be completely corrected by scaling the profile to the correct integral dose. In contrast to that, the profile resulting from the sampled matrix shows little noise (about 3% of the maximum dose value), but by averaging over a small area the dose is correctly predicted even in the outer regions.

Using the dose dependent probability for sampling has an additional potential for reducing the matrix size: Because all sampled values at a fixed radiological depth have the same value, it is no longer necessary to store every dose value separately; instead, this value can be stored only once, together with all its associated voxel indices. This technique is not implemented yet.

The optimization for the three different clinical cases shows that the optimal setting for the dose-dependent sampled matrix D_{ij} is $r_0=25$ mm. With this setting, the treatment plan is clinically equivalent to the respective plan achieved by the fully filled matrix in all three tested cases. As criteria the dose volume histograms and the isodoses were used. The head and neck case was the most challenging due to the tumor size and the problematic locations of the organs at risk; the other two cases could be optimized with even smaller sampled matrices (i.e., D_{ij} matrices with $r_0=20$ mm) without degrading the plan. The head and neck case is presented in detail and the results of the three different (full, sampled and cutoff) optimization methods are compared. Again, for the cutoff technique we used a cutoff value of 45 mm for a comparable matrix size. Table 3.1 lists the key data for the three methods. The resulting optimization time of 45 min is clearly not satisfactory for routine clinical use. Sampling the matrix reduces the memory requirements and speeds up the calculation by a factor of approx. 3. The quality of the final treatment plan is not affected, see also the DVH-comparison in fig. 3.5a.

The maximum difference observed between the two DHVs is a dose deviation of 2.2% of the mean dose at the 80% volume level of the spinal cord. In contrast



Figure 3.5: Dose-volume-histograms for the head and neck case, extracted from dose distributions optimized (a) by a fully filled and a sampled matrix and (b) by a fully filled and a cutoff matrix.

	Full	Sampled	Cutoff
		$(r_o = 25 \mathrm{mm})$	(at $r = 45$ mm)
Number of D_{ij} elements	~110,000,000	~35,000,000	~37,000,000
Total memory used (MB)	~860	~ 380	~ 400
Optimization time (min)	~ 45	~ 15	~ 15
Mean dose (Gy) in:			
Target	60.0	60.0	60.0
Boost	61.9	62.0	62.3
Spinal Cord	22.5	22.4	24.0
Parotid	18.9	19.0	18.3
Brainstem	12.9	12.9	13.1

Table 3.1: Comparison of the optimizations for the head and neck case by different dose calculation matrices. All computations were done on a Compaq Alpha workstation 600 MHz.

to this plan, the simple cutoff approach results in a significantly degraded plan, with a higher dose to the spinal cord. This is also shown in the DVH-comparisons, see fig. 3.5b. Here the maximum difference, again at the 80% volume level of the spinal cord, is 16.4%.



Figure 3.6: Transverse CT slices of the head and case. The outlines of the horseshoe formed target and the spinal cord are shown as black lines. The isodoses of the dose distributions optimized (a) by a full, (b) by a sampled and (c) by a cutoff matrix are labelled with the absolute dose values in Gy. Additionally, the directions of the 7 beams are shown.

Fig. 3.6 shows a transverse slice through the neck region, showing the outlines of the target and the spinal cord, together with the isodoses and the 7 beam directions. The 30 Gy isodose inside the spinal cord is almost identical for the full
and the sampled matrix optimizations, but is located more centrally for the cutoff matrix optimization. Because the dose level rises from inner to outer portions of the spinal cord, this is another illustration of the higher dose deposition in this area.

The degrading of a treatment plan optimized by a cutoff matrix is due to the fact that the optimization algorithm cannot consider scatter contributions to an organ at risk if the central ray of a pencil beam is further away than the cutoff distance; so the algorithm tends to over amplify such a pencil beam to gain the required dose in the target volume.

Because the dose distribution is given by eq. 3.1, it is clear that a dose calculation matrix D_{ij} with 1/3 of the original size results in a three-fold speed gain of the dose calculation. However, it has to be noted that also the number of iterations determines the overall length of the optimization process. Using sampled dose calculation matrices does not increase the number of iterations, although there are small statistical fluctuations which could have slowed down the convergence. This slow down does not occur because in the iterative update of the fluence matrix many voxels are projected to each bixel, averaging the noise out (see also Fig. 3.4b).

The largest deviation between the DVHs for the different techniques is observed in the head and neck case. Also in the other two cases (not shown here) the sampling leads to better or at least equivalent results compared to the cutoff method (with no differences to full-matrix optimization in the DVHs and isodoses for $r_0 = 25$ mm). It should also be noted that the sampling technique shows excellent agreement with the "exact" values for both relative and absolute dose distributions. The cutoff approach does not calculate the absolute dose (i.e., the output factor) well.

The following calculations shown in this thesis were performed with the optimal setting ($r_0=25$ mm for a 1 cm² pencil beam) found in this investigation. In future, tests on more clinical cases should be performed to ensure that this parameter setting is always sufficient to gain an optimal treatment plan. The results obtained so far, especially for the complicated head and neck case which is a real challenge for every inverse planning system, indicate that this is likely to be the case. Compared to measurements, the currently implemented dose calculation algorithm achieves the same accuracy as the pencil beam algorithm by Bortfeld et al. (1993), which is 2-3% for the wide majority of clinical treatment sites. In some cases with complex tissue inhomogeneities (e.g., in the lung and the nasopharynx), the error can be higher. For these cases, more sophisticated dose calculation algorithms like superposition or Monte Carlo methods are preferable. These slow calculations have to be performed only once to build the matrix D_{ij} while the following optimization loop remains unchanged. Therefore it will be more time consuming to build the matrix (for the implemented pencil beam algorithm the build time is about 8 min on a Alpha 600 MHz workstation), but the time needed for the optimization process itself will not increase. The sampling technique presented in this chapter is applicable wherever the dose calculation is performed through a precalculated matrix holding the dose contributions from every bixel to every voxel, independently from the algorithm used for the generation of the matrix.

Chapter 4

The equivalent uniform dose: Models and parameters

Introduced in section 2.2, the concept of equivalent uniform dose (EUD) is a central component of the multicriteria optimization project. Each treatment plan is represented by a vector containing the EUD values of each structure. The definition of Pareto optimality is based on the EUD values, and the EUD is used in the interactive planning session to navigate through the database of treatment plans.

This chapter presents the existing "generalized mean" EUD model and develops the new "max & mean" EUD model. Their parameters are determined by a fit to published clinical tolerance doses, and their predictions for different treated volume fractions are compared and discussed.

4.1 Introduction

Even with the most sophisticated radiotherapy techniques it is impossible to deliver a perfect homogeneous dose to the target volume and no dose to healthy normal structures around it. The dose in each structure will always have a more or less inhomogenous distribution, given by a vector \mathbf{d} containing the dose values of the voxels. Those vectors are not well suited for plan optimization: If there are two plans where one element of \mathbf{d} is higher in the first plan and another element of \mathbf{d} is higher in the second plan, it is hard to decide which plan is

better in a clinical sense. Therefore it is helpful to reduce the complex threedimensional dose distribution to a clinical more meaningful scalar value that allows the ranking of different distributions. One attempt to accomplish this is based on the "effective dose" (Brahme 1994) or "equivalent uniform dose" (EUD) (Niemierko 1997, Niemierko 1999). It is defined as the homogenous dose in an organ that has the same clinical effect as the given, arbitrary inhomogeneous dose distribution. In organs at risk the effect of different spatial dose distributions strongly depends on the organ type. Some organs tolerate very high dose values in small sub-volumes, if the rest of the organ is spared. These organs have a parallel structure: the function is preserved even if a certain fraction of the organ is destroyed. Typical examples are the lung (Kwa et al. 1998a, Theuws et al. 1998), the parotid glands (Eisbruch et al. 1999), the kidneys and the liver. In the other category of organs, high doses are harmful even if they are limited to small volumes - the organ structure is serial. Typical examples for this are the spinal cord and the bowel. The dependency of the tolerance dose on the treated volume is called "volume effect". For different clinical endpoints, there can be different EUD values for the same dose distribution inside an organ, e.g. for pneumonitis and fibrosis in the lung.

The clinical effect of tumor control in the target volume depends on the killing of all clonogen tumor cells. Underdosage even in a small area of the target volume might lead to the survival of some cancer cells and failure of local tumor control. Therefore the EUD in the tumor target volume will mainly depend on the minimum rather than the maximum dose values.

It should be noted that the concept of equivalent uniform dose does *not* tell anything about the *probability* of a clinical effect. The EUD is still a dose, usually given in Gray (Gy). The calculation of the tumor control probability (TCP) and the normal tissue complication probability (NTCP) requires further biological modelling. However, the EUD can serve as an input parameter for such models. For NTCP models, this is equivalent to the 'DVH reduction technique' by Kutcher et al. (1991). Typical TCP/NTCP models have a sigmoidal shape when plotted as a function of EUD.

There are two basic approaches for modelling equivalent uniform dose and TCP/NTCP: the mechanistic and the phenomenological approach. The first one

is based on biological considerations, e.g. cell kill statistics. The latter merely tries to find analytical functions that reproduce given clinical data. The mechanistic approach is inherently more satisfying, but struggles with the complexity of biology. E.g., the mechanistic NTCP model for parallel organs by Jackson et al. (1993) produces an unrealistic NTCP curve that is almost a step function. Only with additional parameters identified with inter-individual heterogeneity of functional reserves and radiosensitivities, the slope becomes shallower. Unfortunately, often these parameters are unknown and cannot be determined independently from the model. At the present time the phenomenological approach seems to be more appropriate for practical use.

One drawback of all existing EUD models to date is that the spatial location of the dose is not taken into account. For an organ at risk with complex substructures it probably has a big impact on the clinical effect where exactly a high dose area is located. Taken the heart as an example, a high dose area in the ventricle with high blood flow should do no harm to the organ, whereas the same dose located at one of the valves might cause complications. Future EUD models might consider the localization of the dose values inside an organ, e.g. by assigning weights to each voxel. At the present day the data basis, i.e. information about 3D dose distributions and their clinical effects, is too small for such sophisticated models. It can be expected that with the broad introduction of 3D planning systems and computerized documentation more clinical data will become available, and the validity of EUD and TCP/NTCP models will improve (Trott 1996).

4.2 The generalized mean model

This section introduces the EUD model that is most widely used in radiotherapy research. It was formulated by Niemierko (1999) in the form

$$\operatorname{EUD}(\mathbf{d}) = \left(\frac{1}{N}\sum_{i=1}^{N} d_{i}^{a}\right)^{1/a}.$$
(4.1)

N is the number of voxels in the structure of interest, $\mathbf{d} = (d_1, .., d_N)$ the dose distribution inside the structure, and a is an tissue-specific parameter. Originally Niemierko formulated the EUD for tumor volumes using a linear-quadratic cell survival model (1997). Eq. 4.1 is valid for both tumors (without considering fractionation) and normal tissues. While it holds its original meaning for tumors (Niemierko 1997), it is a phenomenological model for organs at risk. It can be derived from the power-law relationship that describes the dose-volume effect (Lyman et al. 1987, Mohan et al. 1992, Kwa et al. 1998b):

$$TD(\nu) = TD(1)/\nu^n.$$
(4.2)

Here, TD(1) is the tolerance dose if the whole organ is treated uniformly, and TD(ν) is the tolerance dose if the fraction ν is treated and $(1 - \nu)$ gets no dose. The magnitude of the dose-volume effect is described by the parameter n, whereby n = 1/a.

The EUD as formulated in eq. 4.1 is the "generalized mean" of the dose values. The mathematical properties of the generalized mean are described in the Handbook of Mathematical Functions by Abramowitz and Stegun (1968). Some important properties are:

- If a is negative infinity, the EUD equals the minimum dose.
- If a is 1, the EUD equals the mean dose.
- If a is positive infinity, the EUD equals the maximum dose.
- It is convex (see next chapter for details).

The influence of the parameter a on the EUD is visualized in Fig. 4.1 (taken from Wu et al. 2002). Usually target type structures will have a large negative tissue parameter a, so the EUD is mainly determined by the lowest dose values. Parallel organized organs at risk like the lung have a parameter near 1 and an EUD near the mean dose. The parameter for serial organs like the spinal cord is large positive, so the highest dose values have the biggest impact on the EUD.

4.3 The max & mean model

In this section an even simpler approach to describe the EUD in organs at risk is formulated (see also Thieke et al. 2002b). It is based on the work by Dale and



Figure 4.1: The left panel shows the dose volume histogram of the dose distribution in some structure of interest. On the right panel, the EUD of this dose distribution is plotted as a function of the tissue parameter a. Taken from Wu et al. (2002)

Olsen (1997) who reported that the normal tissue complication probabilities of purely serial and purely parallel organs at risk correspond to the maximum dose and the mean dose in the organ, respectively. For organs with a mixed serial and parallel structure, Dale and Olsen (1997) suggested to use a linear combination of the maximum and the mean dose. In the following, this is called the max & mean approach. The organ-specific parameters for this approach are determined, and the hypothesis is investigated and tested.

The following linear max & mean model is used to represent the EUD:

$$EUD = \alpha D_{\max} + (1 - \alpha)\overline{D}.$$
(4.3)

 D_{max} is the maximum and \overline{D} the mean dose of the radiation delivered to the organ. α is an organ-specific parameter ranging from 0 to 1.

The values of α and EUD were calculated by a fit to the Emami tables (Emami et al. 1991). In that paper tolerance doses (TD) for various organs were presented depending on the irradiated relative volume ν (whereby ν takes values of $\nu_1 = 1/3$, $\nu_2 = 2/3$ and $\nu_3 = 1$) and on two levels of complication rate (5% and 50%, in the following specified as TD5 and TD50) 5 years after treatment. For each complication rate, the associated maximum and mean dose can be determined straightforward: Because the fraction ν gets the dose TD and the fraction $(1 - \nu)$ gets no dose, the maximum dose is given as $D_{\text{max}} = \text{TD}$ and the mean dose is given as $\overline{D} = \nu \cdot \text{TD}$. From these data the organ-specific parameter α can be deduced, which should be independent of the complication rate, and two EUD-values (EUD5 and EUD50), applying to a complication rate of 5% and 50%, respectively. To perform the fit, we rewrite the max & mean model as $\alpha(\overline{D} - D_{\text{max}}) + \text{EUD} = \overline{D}$ and build an overdetermined equation system for each individual organ:

$$\begin{pmatrix} (1-\nu_{1}) & 1 & 0\\ (1-\nu_{2}) & 1 & 0\\ (1-\nu_{3}) & 1 & 0\\ (1-\nu_{1}) & 0 & 1\\ (1-\nu_{2}) & 0 & 1\\ (1-\nu_{3}) & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \alpha\\ \text{EUD5}\\ \text{EUD50} \end{pmatrix} = \begin{pmatrix} \nu_{1} \cdot \text{TD5}(\nu_{1})\\ \nu_{2} \cdot \text{TD5}(\nu_{2})\\ \nu_{3} \cdot \text{TD5}(\nu_{3})\\ \nu_{1} \cdot \text{TD50}(\nu_{1})\\ \nu_{2} \cdot \text{TD50}(\nu_{2})\\ \nu_{3} \cdot \text{TD50}(\nu_{3}) \end{pmatrix}$$
(4.4)

This can be abbreviated as $A \cdot x = c$, and a least-square fit of the three parameters α , EUD5 and EUD50 can be performed by calculating $x = (A^T A)^{-1} A^T c$. As errors the mean quadratic deviation of all meaningful quadratic subsystems (i.e., subsystems with 3 equations) from the least-square fit are used. "Meaningful" are those subsystems which lead to an exact solution for α , EUD5 and EUD50 (e.g. it is not meaningful to select the first three lines of the above equation system because there would be no solution for EUD50).

With regard to the simple data structure provided by Emami et al. (1991), the universal formulation of the generalized mean model,

$$\operatorname{EUD}(\mathbf{d}) = \left(\frac{1}{N}\sum_{i=1}^{N} d_{i}^{a}\right)^{1/a},\tag{4.5}$$

can be simplified to (cf. eq. 4.2, EUD=TD(1))

$$EUD = \nu^{1/a} \cdot TD(\nu), \qquad (4.6)$$

which leads to values for a and EUD in a very similar way as for the max & mean model.

Organ	α	EUD 5 $[Gy]$	EUD 50 $[Gy]$
Larynx (necrosis)	0.83 ± 0.13	68.6 ± 4.1	78.3 ± 4.1
Larynx (edema)	1.00	45.0	
Lung I	0.09 ± 0.09	18.7 ± 2.4	26.0 ± 2.9
Heart	$0.53 {\pm} 0.12$	39.7 ± 2.9	48.2 ± 3.0
Esophagus	$0.90 {\pm} 0.03$	55.7 ± 0.7	67.6 ± 0.9
Stomach	$0.83 {\pm} 0.07$	51.7 ± 2.2	63.5 ± 2.6
Small intestine	$0.80 {\pm} 0.09$	41.7 ± 2.5	53.6 ± 4.3
Colon	$0.75 {\pm} 0.02$	$45.4{\pm}0.7$	54.6 ± 0.8
Liver	$0.50 {\pm} 0.12$	30.8 ± 3.0	38.0 ± 2.9
Kidney I	0.17 ± 0.08	22.4 ± 0.8	28.5 ± 1.0
Bladder	$0.64{\pm}0.19$	67.7 ± 4.5	77.4 ± 2.6
T-M joint mandible	$0.89 {\pm} 0.08$	59.4 ± 2.3	71.0 ± 2.6
Brain	$0.67 {\pm} 0.06$	$45.4{\pm}1.8$	58.7 ± 1.9
Brain stem	$0.75 {\pm} 0.04$	49.5 ± 1.0	
Brachial plexus	$0.96 {\pm} 0.00$	60.1 ± 0.2	74.9 ± 0.2
Ear mid/ext (acute)	1.00	30.0	40.0
Ear mid/ext (chronic)	1.00	55.0	65.0
Parotid I and II	1.00	32.0	46.0

Table 4.1: Fitting the max & mean model to the data by Emami et al. (1991).

4.4 Results

Table 4.1 shows the result of the fits (the particular clinical endpoints are only noted for organs with more than one endpoint, the other ones can be found in (Emami et al. 1991). For the larynx (endpoint edema) and the brain stem only data for a complication rate of 5% are provided. Wherever no error is listed in table 4.1, only two $TD(\nu)$ -values are given by Emami et al. (1991), making the fit mathematical exact.

Once the values of α and EUD respectively a and EUD have been fitted, the tolerance dose $\text{TD}(\nu)$ of an arbitrary irradiated fraction ν can be calculated. To compare the generalized mean model with the max & mean model, the tolerance dose $\text{TD}(\nu)$ is plotted as a function of the irradiated fraction ν . Fig. 4.2 show these plots for the cases of the lung (α =0.09, a=1.18), the heart (α =0.53, a=2.94) and the esophagus (α =0.90, a=16.67), representing different organ structures from "mostly parallel" to "mostly serial". Both models fit the data within a tolerance of about 10%.



Figure 4.2: Tolerance doses in dependency on the irradiated volume fraction. Shown are the discrete values given by Emami et al. (1991) and the comparison of the generalized mean and the max & mean EUD models. (a) Lung, (b) Heart and (c) Esophagus.

4.5 Discussion

The results show that the max & mean model can be fitted to the Emami tables with the same accuracy as the standard generalized mean model. A single value of α suffices to characterize most organs at different tolerance dose levels. Whether the model can also be fitted to real clinical data remains to be seen when more data become available. It has to be noted that all error ranges shown in this chapter are the fit errors with respect to the tables by Emami et al. (1991). The errors of the Emami data themselves, which for some organs are highly speculative, are not considered.

Especially from figure 4.2c it is obvious that the max & mean model differs from the generalized mean model in the case of small treated sub-volumes. Mathematically this is clear because the power law expression of the generalized mean model has a pole at $\nu = 0$, i.e. according to this model small sub-volumes can be treated with extremely high doses. This is not the case for the max & mean model, which is quite sensitive to small hot spots. On the one side this can be advantageous: If the model is used for plan evaluation in inverse treatment planning, overdosages in small sub-volumes of a serial organ can be effectively reduced; in some cases the use of the generalized mean model or quadratic deviations from a subscribed physical dose penalize such overdosages not strong enough. On the other hand, an organ at risk lying directly next to the tumor target volume has to get a high dose at the border in order to achieve the curative dose in the tumor volume. High doses in very few organ voxels should therefore be made possible.

So a reasonable way to deal with the characteristic of the max & mean model is the appropriate determination of D_{max} : It should be identified with the highest mean dose over a certain sample volume, e.g. about 1.5 cm³, rather than the highest dose value of one single point. This definition follows ICRU 50 (1993) and ensures that the dose values in areas with a steep dose gradient (i.e. at the border of organs at risk lying directly next to the target volume) are not overestimated.

Again, also for small treated sub-volumes more clinical data are desirable for a better evaluation of the different models.

A benefit of the max & mean approach is that it can be applied to retrospective data even if only the maximum and the mean doses are known. However, for organs with a steep dose gradient it is important to know the exact definition of the maximum dose. It remains to be investigated whether a similar approach as the max & mean approach can also be used to characterize the minimum dose constraint in the tumor target volume. Obviously, since here the minimum dose is more relevant than the maximum dose, one should rather use a "min & mean" approach.

One main advantage of the max & mean EUD concept is its linearity. It allows the formulation of the multicriteria optimization problem in form of a linear program. The EUD based on the max & mean model can be used in such a program as maximum constraint for organs at risk. As minimum constraint for the target volume, the standard minimum physical dose constraint can be used. The linear program has the following setting:

Subject to the constraints

$$EUD(D_k) \le U_k(1+t_k), \quad k=1,\ldots,K \quad \text{``risk constraints''}$$
(4.7)

and

$$\operatorname{Min}(D_{\mathrm{T}}) \ge L(1-t)$$
 "target minimum constraint", (4.8)

independently minimize the objectives

$$t_k \to \min, \quad k=1,\ldots,K; \quad t \to \min.$$
 (4.9)

Here, the numbers U_k are ideal tolerance doses for organs at risk while L is an ideal curative dose for the target volume. The target volume is additionally considered as an organ at risk to ensure dose homogeneity.

The max & mean model and the generalized mean model can be fitted within a similar error range (about 10%) to the presently available (Emami) data. At this time, it is not clear which model describes an inhomogenous dose distribution with a higher clinical relevance. Therefore it seems appropriate to implement both models into the multicriteria optimization project.

The linear max & mean model has good numerical properties and can be solved in standard linear programming tools that are commercially available. The collaboration partners of the multicriteria optimization project at the Fraunhofer Institute for Industrial Mathematics are working on the generation of Pareto optimal solutions using the max & mean model and linear solver tools. Because of the power-law formula, the generalized mean EUD model is in general not linear. Commercially available linear solvers cannot be used for optimizing plans using this EUD model. In the next chapter a new algorithm that solves this problem is developed and tested.

Chapter 5

The equivalent uniform dose as an optimization constraint

Optimization algorithms in inverse radiotherapy planning need information about the desired dose distribution. Usually the planner defines physical dose constraints for each structure of the treatment plan, either in form of minimum and maximum doses or as dose-volume constraints. In the multicriteria optimization project the equivalent uniform dose (EUD) was chosen as criterion for each structure in the treatment plan. In consequence it should be the EUD instead of physical dose that is used as an optimization constraint in the calculation of treatment plans.

In this chapter, a method is presented to consider the generalized mean EUD model as an optimization constraint by using projections onto convex sets (POCS). The actual dose distribution is repeatedly projected onto the set of distributions that fulfil the EUD constraint, leading to voxel-based physical dose constraints. The algorithm is easy to integrate into the existing inverse planning systems. It can also be used in existing (weighted and constrained) inverse optimization setups as well and not exclusively in multicriteria optimization. It is possible to choose between physical and EUD constraints separately for each structure, allowing the planner to make use of the EUD concept in conventional inverse planning and carefully test this new concept in radiotherapy optimization.

The following sections (see also Thieke et al. 2002c) are therefore valid for both the existing inverse planning concepts and the multicriteria optimization project.

5.1 Introduction

In current inverse treatment planning programs for intensity-modulated radiotherapy, the planner defines certain constraints for the desired dose distribution, and in an optimization process the computer tries to find the treatment setup that matches the constraints as closely as possible. Common constraints for the optimization are minimum and maximum doses for target structures and maximum doses for normal structures, each in conjunction with a penalty factor (Bortfeld et al. 1999).

These physical dose constraints allow the implementation of fast converging and robust gradient optimization algorithms. However, finding the best parameter setting is non-intuitive. Because often the dose constraints cannot be fulfilled for every single voxel, the penalty factors have a great impact on the optimization result. In these cases, the prescribed maximum dose loses its original meaning of a real constraint and is used as a mere steering parameter for the optimization outcome. Balancing the maximum/minimum constraints and their penalties is a trial and error process with several optimization runs.

For some organs at risk, a pure maximum dose constraint is not meaningful at all. As presented in the previous chapter, parallel organs like liver, lung and kidneys show a distinct volume effect, i.e. the tolerance dose for small irradiated volume fractions is much higher compared to irradiation of large fractions. To consider volume effects in a physical dose constrained optimization, the planner can define dose volume histogram (DVH) constraints, i.e. instead of a maximum dose, the constraints are formulated in the form "not more than x % of the organ are allowed to receive more than y Gy", see e.g. Bortfeld et al. 1997. Usually several DVH constraints have to be defined for one organ. This increases the number of parameters to be defined by the treatment planner, and it may exclude solutions that keep organs at risk inside the tolerance, but that have a different shape of the DVH.

The concept of EUD is probably of higher clinical relevance than single physical dose values and is therefore a promising concept for overcoming the difficulties of physical dose constrained optimization. Wu et al. (2002) used the generalized mean based EUD for inverse treatment planning as a parameter in a sigmoid dose-effect curve that resembles the basic shape of TCP/NTCP models. In this chapter, the generalized mean based EUD is directly implemented as an optimization constraint without assuming any dose-effect-relationships. In the following, the mathematical aspects of the implementation are described, the differences between physical dose constrained and EUD constrained optimization are shown, and the clinical application is discussed.

5.2 Material and Methods

The calculation platform for the investigation is the same as in chapter 2. As an initial step prior to the optimization loop, the dose contribution of every beam element (bixel) j to every voxel i is stored for unit fluence in a matrix D_{ij} . For arbitrary beam weights $\mathbf{w} = (w_1, ..., w_M)$ the dose $\mathbf{d}_T = (d_1, ..., d_{N_T})$ can then be calculated through

$$d_i = \sum_{j=1}^{M} D_{ij} w_j \qquad i = 1, .., N_T.$$
(5.1)

M is the number of bixels, and N_T the total number of voxels.

The planning system uses a Quasi-Newton gradient technique to optimize the fluence maps. The objective function is defined as

$$F = \sum_{i=1}^{N_T} s_i \left(d_i - d_i^{\text{pres}} \right)^2.$$
 (5.2)

 d_i^{pres} and s_i are the desired dose for the voxel *i* and the associated penalty factor, based on the planner's prescription and the current dose. For minimum and maximum dose constraints, d_i^{pres} and s_i are given by

$$\left(d_i^{\text{pres}}, s_i \right) = \begin{cases} \left(d_i^{\min}, s_i^{\min} \right) & \text{if } d_i < d_i^{\min} \\ \left(d_i, 0 \right) & \text{if } d_i^{\min} \le d_i \le d_i^{\max} \\ \left(d_i^{\max}, s_i^{\max} \right) & \text{if } d_i > d_i^{\max} \end{cases}$$
(5.3)

 (d_i^{\min}, s_i^{\min}) and (d_i^{\max}, s_i^{\max}) are the minimum and maximum physical dose constraints and penalty factors for voxel *i*. Note that the system internally works with individual settings for each voxel. The planner usually defines the dose constraints and penalties only separately for each structure (targets/organs at risk), which are then assigned to all voxels belonging to the particular structure. Naturally, for organs at risk the minimum dose constraint is set to 0.

DVH constraints are considered by an extension of the above concept, whereby up to 5 DVH constraint points per structure are transformed into adequately chosen values of d_i^{pres} (Bortfeld et al. 1997).

The iterative update of the beam weights from iteration (t) to iteration (t+1) is given by

$$w_{j}^{(t+1)} = \left[w_{j}^{(t)} - \alpha \left(\frac{\mathrm{d}F}{\mathrm{d}w_{j}} / \frac{\mathrm{d}^{2}F}{\mathrm{d}w_{j}^{2}} \right) \right]_{+} = \left[w_{j}^{(t)} - \alpha \frac{\sum_{i=1}^{N} 2s_{i} \left(d_{i} - d_{i}^{\mathrm{pres}} \right) D_{ij}}{\sum_{i=1}^{N} 2s_{i} D_{ij}^{2}} \right]_{+}$$
(5.4)

 α is a damping dactor to ensure convergence, and because the fluence has to be non-negative, the operator

$$[x]_{+} = \begin{cases} x & \text{if } \mathbf{x} \ge 0\\ 0 & \text{if } x < 0 \end{cases}$$
(5.5)

is used.

The handling of physical dose constraints is also illustrated in Fig. 5.1.



Figure 5.1: Illustration of maximum and DVH constraints in inverse planning. Shown is the actual dose at some point during the optimization in an organ at risk in form of a dose-volume-histogram. Both the maximum dose constraint and the DVH constraint for this organ are violated in this example, and so the optimization algorithm tries to find the dose distribution whose DVH avoids the hatched areas. All other parts of the organ dose can remain unchanged.

It shows the DVH of a dose distribution in an organ at risk at some arbitrary point of the optimization. One DVH constraint and a maximum physical dose constraint are defined for this organ, and both are violated by the actual dose distribution. The planning system then generates the arrays d^{pres} and s^{pres} in a way that only those voxels violating the constraints contribute to the objective function. For all other voxels the new prescribed dose equals the actual dose. The actual dose distribution is modified as little as possible, but also as much as necessary to match all constraints.

This concept has proven to result in a fast and stable optimization. It was the motivation to implement EUD constraints into the system in a way that is compatible with the existing concept and that keeps most parts of the established algorithm unchanged. The generalized mean EUD model as introduced in chapter 4 (cf. eq. 4.1) is defined as

$$\operatorname{EUD}(\mathbf{d}) = \left(\frac{1}{N}\sum_{i=1}^{N} d_{i}^{a}\right)^{1/a}.$$
(5.6)

Because the EUD is calculated separately for each structure, in the following the dose vector **d** is restricted to those voxels belonging to the particular structure and has N elements, $\mathbf{d} = (d_1, .., d_N)$.

Let us now assume that at any stage during the optimization we have a dose distribution **d** inside the organ that violates a predefined EUD constraint, e.g. $EUD(\mathbf{d}) > EUD^{\max}$ for an organ at risk. In analogy to the handling of physical dose constraints described above, we now want to find a new dose distribution \mathbf{d}' that fulfils the constraint, i.e. $EUD(\mathbf{d}') \leq EUD^{\max}$, and at the same time is as close to **d** as possible. The elements of \mathbf{d}' , $(d'_1, ..., d'_N)$, can then be used as physical dose constraints in the array \mathbf{d}^{pres} , see eq. (5.2).

To find \mathbf{d}' , we make use of a mathematical property of the EUD formula, its convexity for a < 0 and $a \ge 1$ (Choi 2002). Convexity means that for two given dose distributions within the constraint, any convex combination of these two is also within the constraint:

$$\alpha \text{EUD}(\mathbf{d}_1) + (1 - \alpha) \text{EUD}(\mathbf{d}_2) \le \text{EUD}^{\max}$$
(5.7)

for EUD
$$(\mathbf{d}_1)$$
, EUD $(\mathbf{d}_2) \le \text{EUD}^{\max}$, $0 \le \alpha \le 1, a \ge 1$ (5.8)

and

$$\alpha \text{EUD}\left(\mathbf{d}_{1}\right) + (1 - \alpha) \text{EUD}\left(\mathbf{d}_{2}\right) \ge \text{EUD}^{\text{IIIII}}$$

$$(5.9)$$

for EUD
$$(\mathbf{d}_1)$$
, EUD $(\mathbf{d}_2) \ge \text{EUD}^{\min}, 0 \le \alpha \le 1, a < 0.$ (5.10)

In other words, all dose distributions that satisfy a given EUD constraint form a convex set inside the N-dimensional dose space. This property is essential to make the projection from \mathbf{d} to \mathbf{d}' unique. It is called a projection onto a convex set (POCS) (Bregman 1965, Gubin et al. 1967), a technique known from other applications, e.g. in image restoration (Youla et al. 1982). POCS has also already been used in radiotherapy planning (Lee et al. 1997), especially as an alternative method for considering DVH constraints (Cho et al. 1998). Fig. 5.2 illustrates the principle.



Figure 5.2: The principle of Projection onto convex sets (POCS).

For reasons of clarity, an organ with only two voxels (N = 2) is shown, so that every dose distribution $\mathbf{d} = (d_1, d_2)$ can be represented by a point on the plane. We look at an organ at risk with an EUD parameter of a = 2 and a prescribed maximum constraint EUD^{pres}. All dose distributions fulfilling the constraint form a convex set that includes $\mathbf{d} = (0, 0)$ and has a border forming a circle segment from $\mathbf{d} = (\sqrt{2} \cdot \text{EUD}^{\text{pres}}, 0)$ to $\mathbf{d} = (0, \sqrt{2} \cdot \text{EUD}^{\text{pres}})$. For other organ parameters a, the border would have a different shape, but would be still convex. Also shown in Fig. 5.2 is a dose distribution \mathbf{d} that violates the constraint (and is therefore outside the convex set) and its projection to the nearest distribution \mathbf{d}' of the convex set. It is clear that \mathbf{d}' will fulfill the EUD constraint exactly and does not go beyond that, i.e. we will have EUD (\mathbf{d}') = EUD^{pres} instead of the inequality used so far. The same applies for minimum EUD constraints, so the following calculation is valid for both maximum and minimum constraints.

To calculate \mathbf{d}' , we have to find the minimum of

$$f(\mathbf{d}') = \sum_{i=1}^{N} (d_i - d'_i)^2$$
(5.11)

under the constraint $\text{EUD}(\mathbf{d}') = \text{EUD}^{\text{pres}}$. This is a classical setup for the use of Lagrange multipliers. We define a Lagrange function

$$L(\mathbf{d}') = f(\mathbf{d}') + \lambda \left(\text{EUD}^{\text{pres}} - \text{EUD}(\mathbf{d}') \right)$$
(5.12)

and set the gradient to zero:

$$\left(\frac{\partial L}{\partial \mathbf{d}'}\right)_{j} = -2\left(d_{j} - d_{j}'\right) - \frac{1}{N}\lambda d_{j}'^{a-1} \left(\frac{1}{N}\sum_{i=1}^{N}d_{i}'^{a}\right)^{\frac{1}{a}-1} = 0$$
(5.13)

$$\Rightarrow \quad \frac{d_j - d'_j}{d'_j^{a-1}} = -\frac{\lambda}{2N} \left(\frac{1}{N} \sum_{i=1}^N {d'_i^a}\right)^{\frac{1}{a}-1} \quad j = 1, .., N \tag{5.14}$$

The explicit value of the Lagrange multiplier λ is unknown, but it can be seen from eq. (5.14) that for every voxel index j there is a constant expression for d'_{j} :

$$\frac{d_j - d'_j}{d'_j^{a-1}} = \text{const} \quad j = 1, .., N.$$
(5.15)

The value of const that is exact for a = 1 and a = 2 and approximately correct for all other values of a is given by

$$\frac{d_j - d'_j}{d'_j^{a-1}} \approx \frac{\text{EUD}\left(\mathbf{d}\right) - \text{EUD}^{\text{pres}}}{\left(\text{EUD}^{\text{pres}}\right)^{(a-1)}}.$$
(5.16)

Eq. (5.16) is an implicit definition for \mathbf{d}' . In practice, this is as good as an explicit definition, because each element d'_j can be calculated in a short and fast iterative sub-calculation. It is interesting to note that for a = 1 the dose is just shifted by const = $(EUD(\mathbf{d}) - EUD^{\text{pres}})$, and for a = 2 the relative change of the dose is constant.

When projecting to a lower EUD, for every j it is

$$d'_{j} \le d_{j} \quad j = 1, .., N,$$
 (5.17)

and when projecting to a higher EUD, used for target structures in case the minimal EUD constraint is violated, we have

$$d'_{i} \ge d_{i} \quad j = 1, .., N.$$
 (5.18)

If only a minimum EUD constraint is used for target structures, eq. (5.18) would lead to very high doses in some target voxels because they would never be actively reduced. To assure target dose homogenity, we therefore consider the target also as an organ at risk and prescribe a maximum EUD. There has to be a separate target parameter $a \ge 1$ for the maximum EUD constraint to ensure the convexity. The use of both minimum and maximum EUD constraints for target structures implies that there might be two projections directly after one another to obtain \mathbf{d}' : First, to a higher EUD in case the minimum EUD constraint to to the maximum constraint.

In summary, the EUD constraint defined by the planner is internally transformed into physical dose constraints individual for each voxel. In practice, this is a simple replacement of eq. (5.3) by eq. (5.16). The objective function, its gradients and all other parts of the optimization algorithm can be retained without change. It is even possible to use EUD constraints only for some organs and physical dose constraints for the others. This is illustrated in the flowchart in Fig. 5.3. The well-established optimization program could be retained without change, and the feature of EUD-constrained optimization is just a "plug-in" (the part of the figure with a dotted background) into the routine that builds the array d^{pres} . From Fig. 5.3 one can also see how mixing of different constraint types is implemented: Because the array d^{pres} is calculated in a sub-loop organ by organ, the planner can define separately for each organ either physical max/min, DVH or EUD constraints.



Figure 5.3: Optimization loop of the inverse planning program used for this paper. In each iteration, a dose distribution \mathbf{d}^{pres} is generated, based on the current dose and individual for each voxel. \mathbf{d}^{pres} is generated by processing organ after organ, and so it is possible to chose different constraint types separately for each of them. Because the POCS method leads to voxel-based physical dose constraints, EUD constraints could be implemented by simply adding a branch to the generation of \mathbf{d}^{pres} (marked by the dotted background in the figure).

5.3 Results

First, we will look at projections for organs at risk that violate a maximum EUD constraint. As it can be seen from eq. (5.16), the projection from the current dose distribution to the one that satisfies the EUD constraint depends on the tissue parameter a. For serial organs, the EUD is mostly determined by the highest voxel doses, so it would have little effect to change the parts of the organ that already receive a low dose. In more parallel organs, the EUD is nearer the mean dose, and the voxels are affected more evenly by the projection.



Figure 5.4: Examples of projections to a lower EUD for different organ parameters 'a'. For a=1.0 (parallel organ), the DVH is shifted to the left. For a=2.0, the DVH is scaled along the dose axis, and for a=7.4 (the parameter for the spinal cord), higher doses are even more effected and lower dose parts are not changed at all. The extreme case a= ∞ means the equivalence of EUD and physical dose maximum and would lead to a projection shown in fig. 5.1.

Fig. 5.4 shows the DVHs of an arbitrary dose distribution in an organ at risk at some point during the optimization, and its projections assuming three different parameters a. The EUD constraint for this organ was set to 33 Gy. Depending on a, the EUD of the current dose distribution is 50.8 Gy (a = 7.4, which is the parameter for the spinal cord, a mostly serial organized structure), 45.4 Gy (a = 2.0) and 40.9 Gy (a = 1.0, representing a purely parallel organized)organ). For all three parameters the constraint of 33 Gy was violated, and for each parameter the algorithm has to project to a new dose distribution. The resulting projections are also shown in Fig. 5.4 and have an EUD of 32.4 Gy (a = 7.4), 32.9 Gy (a = 2.0) and 34.0 Gy (a = 1.0). One can see that for a = 7.4and a = 2.0 the EUD of the projected dose distribution meets the constraint and is very close to it. Only for a = 1.0, the EUD of the projection is slightly higher than the constraint. This is because the dose has to be non-negative, and therefore shifting the dose is not completely possible in voxels with a dose $d_i < (EUD(d) - EUD^{\text{pres}})$. But this effect is small and becomes even smaller after each iteration when the constraint is reached more closely, and can be neglected.

As stated in the previous chapter, for targets the situation is slightly more complicated because there are both minimum and maximum constraints. Two parameters are required, a negative a_{-} for the minimum EUD constraint and a positive a_{+} for the maximum EUD constraint. It is possible that the current dose inside the target violates both constraints, or that the maximum constraint is violated after the projection to the minimum EUD. In these cases two projections directly after one another are necessary. Fig. 5.5 shows an exampe for a target volume with the constraints EUD^{min} = 60 Gy ($a_{-} = -8$) and EUD^{max} = 61 Gy ($a_{+} = 10$). First, the minimum EUD constraint is reached by the first projection. Then, the second projection ensures that also the maximum EUD constraint is met.

In the following, we look at optimization results for a clinical case using the common physical constraints and the new EUD constraints. The case we looked at was an extensive head and neck tumor with a complex shaped target volume, including a boost volume, and located very close to organs at risk: brainstem, spinal cord and parotid gland (the same as in chapter 3, cf. fig. 3.3). Table 5.1 shows parameter settings and optimization results. Treatment plan A was calcu-



Figure 5.5: For target structures, there is a minimum EUD constraint (associated with a negative parameter 'a') and, to ensure homogeneity, also a maximum EUD constraint (associated with a positive parameter 'a'). This can lead to two projections in one iteration to find a prescribed dose distribution that fulfills both constraints. In this example the constraints were set to $EUD^{min}=60$ Gy and $EUD^{max}=61$ Gy.

	Plan A	Plan B	Plan C
	(physical)	(physical/EUD)	(EUD)
Optimization			
constraints (Gy):			
Target $(a=-8)$	Dmin=60	Dmin=60	EUDmin=60
$(a_{+}=10)$	Dmax=60	Dmax=60	EUDmax=61
Boost $(a=-8)$	Dmin=62	Dmin=62	EUDmin=62
$(a_{+}=10)$	Dmax=62	Dmax=62	EUDmax=63
Spinal cord $(a=7.4)$	Dmax=33	EUDmax=25.5	EUDmax=25.5
Brainstem $(a=4.6)$	Dmax=45	EUDmax=23	EUDmax=23
Parotid gland (a=1)	Dmax=25	EUDmax=13	EUDmax=13
Result EUD (Gy):			
Target $(a=-8)$	58.5	58.5	58.5
$(a_{+}=10)$	60.2	60.5	60.7
Boost $(a=-8)$	61.8	62.1	61.6
$(a_{+}=10)$	62.0	62.3	62.7
Spinal cord $(a=7.4)$	28.8	25.8	25.4
Brainstem $(a=4.6)$	27.7	23.1	22.9
Parotid gland $(a=1)$	17.6	13.5	13.0
Iterations	216	309	500

Table 5.1: Optimization constraints and results for the head and neck case. For plan A, only physical constraints were used, for plan B physical constraints for the target/boost volumes and EUD constraints for the organs at risk were used, and plan C was optimized using EUD constraints for all structures. All plans were normalized to a target EUD of 58.5 Gy.

lated with maximum and minimum physical dose constraints for all structures. In treatment plan B the organs at risk were optimized using EUD constraints while keeping the physical constraints for the target, and plan C is the optimization result using EUD constraints for both organs at risk and the target. The slightly more elaborate generation of the prescribed dose array $\mathbf{d}^{\mathrm{pres}}$ for EUD constraints (compare eqs. (5.3) and (5.16)) is not noticable in terms of computation time: The time needed for one single iteration is the same for physical and for EUD constraints; in this specific case this is approx. 7 sec on a Pentium IV 1.4 GHz workstation. But the total number of iterations needed increases for EUD constrained optimization, cf. Table 5.1. Fig. 5.6 shows both the DVHs of plan A and plan B. The EUD values of all structures in plan A and the target EUDs in plan B are displayed only for informational purposes; they were not used during the optimization. Plan B with EUD constraints shows better sparing of the organs at risk while keeping the dose in the target at the same level as plan A. For the parotid gland (a parallel organ with a = 1), in plan B the maximum dose values are higher, but the mean dose (and therefore the EUD) is lower. For all structures, the EUD constraints are reached very closely. In Fig. 5.7, the DVHs of plan C and, again, plan B is shown. The results are similar, but the homogeneity of the target, especially the boost, in plan C is worse than in plan В.

5.4 Discussion

As can be seen in the results section, the algorithm presented is capable of finding a plan that reaches given EUD constraints closely. It needs more iterations than the physical constrained optimization, but using EUD constraints only for organs at risk, the convergence is still good. Although in the test case the EUD constrained optimization shows better sparing of organs at risk than the physical constrained optimization, this cannot be claimed in general. The results depend directly on the concrete constraint settings. It is also important to note that the EUD of an organ at risk is not optimized beyond the constraint. Without further modifications to the algorithm (which are presented in section 6.2.1), the planner has to define a sufficiently low constraint to minimize the EUD as much as possible.



Figure 5.6: Optimization results for the head and neck case. Plan A was generated using physical maximum/minimum constraints for all structures. For plan B the target and boost constraints were the same as in plan A, but the organs at risk were optimized using maximum EUD constraints.



Figure 5.7: Again optimization results for the head and neck case. Plan B is the same as in figure 5.6, now compared to a plan C where EUD constraints were used for all structures.

The same optimization result as in plan B could have been also reached with a set of DVH constraints. However, these DVH constraints were unknown, and finding the best set would require a lot of trial and error (Wu et al. 2002). The fact that EUD constraints can lead to a better overall result in the organs at risk than simple maximum dose constraints becomes plausible by looking at Fig. 5.1 and Fig. 5.4: Whereas no voxel below the maximum dose constraint is optimized (see Fig. 5.1), the EUD constraint affects, depending on the organ parameter, much more or even all voxels of the organ (see Fig. 5.4).

The question is where the use of EUD constraints is most suitable and how they compare to the physical constraints. From eq. (5.6) one can see that when ais positive infinity (completely serial organized organ), the EUD equals the maximum dose in the structure. In this case, the POCS method projects exactly onto the maximum dose constraint, so the use of "maximum physical dose constraint" and "maximum EUD constraint" is equivalent. In analogy, a set to negative infinity (where the EUD equals the minimum dose) implies the equivalence of a minimum EUD constraint and a pure physical dose minimum constraint. In that sense the pure physical maximum/minimum dose constraints are a subset of EUD constraints, namely for structures with a set to $+/-\infty$.

The differences become bigger and bigger when the organ structure becomes more and more parallel (cf. parotid gland in Fig. 5.6), indicating that a physical maximum dose constraint is no longer sufficient. Especially for these organs, either DVH constraints or EUD constraints should be used. The EUD constraint is easier to define and does not imply a predefined shape of the DVH. The optimization will find the dose distribution that is best adapted to the specific case. As found in the results section, even for organs considered as mostly serial (like the spinal cord) EUD constraints can improve the optimization result.

EUD constraints for target structures using organ parameters as in plan C resulted in more iterations and lower dose homogeneity. So for targets the use of higher a_+ (for the max. EUD constraint) and lower a_- (for the min. EUD constraint) organ parameters up to positive/negative infinity seems more suitable.

The algorithm presented makes it possible to integrate the non-linear generalized mean EUD model into the multicriteria optimization concept. This will be utilized in the next chapter where a complete set of Pareto optimal solutions is generated and stored in a database.

In addition to the implementation in multicriteria optimization, another main advantage of the algorithm presented is that it can be used in the common inverse planning concepts. It is very close to the existing physical optimization and is easy to integrate into any inverse planning system that uses a gradient optimization technique and voxel based dose constraints. Virtually nothing of the existing system has to be changed; the EUD constraint is just an additional feature and is realized by a simple plug-in, cf. Fig. 5.3. For every organ or structure the planner can decide separately what kind of constraint should be used. It is not necessary to change the whole planning paradigm. Inverse planning is a process that requires a lot of experience from the planner, and the introduction of EUD constraints as presented does not render this experience obsolete. Instead, it is possible to make use of the EUD concept in a conservative way by using an EUD constraint only for selected organs, and use physical dose constraints and the knowledge about the correct settings for constraints and penalty factors for all other organs. This ensures a smooth transition from a pure physical dose oriented to a clinical more meaningful optimization. When more and more clinical data become available and the parameters for tissue parameters and tolerance doses become more valid, the planner can make use of it by using the EUD feature for more organs. The usefulness of a smooth transition is also supported by the fact that planning based on TCP/NTCP models was introduced several years ago (Sodertrom et al. 1993, Niemierko et al. 1992) but is still not widely used in clinical practice. The EUD is an intermediate concept between physical doses and TCP/NTCP models. It is still in the dose domain and does not cope with, but on the other hand is also not affected by the uncertainties of dose-effect relationships.

Chapter 6

Practical application of multicriteria optimization in inverse radiotherapy planning

The fundamental algorithms and analyses described in the previous chapters now allow the practical application of multicriteria optimization. This chapter describes the first implementation of a multicriteria optimization scheme in inverse intensity-modulated radiotherapy planning. After presenting a simple clinical test case, some techniques that might become important for more complex cases are introduced as well.

6.1 Clinical test case with 1 target volume and 2 organs at risk

6.1.1 Introduction

The most simple anatomic situation for radiotherapy optimization is one single target volume next to one single organ at risk. However, for such a case there exist also only one single Pareto optimal solution (assumed that there is only one criterion representing the target dose distribution). Different doses to the target volume can be achieved by scaling the whole plan to different dose levels. So, the clinical case investigated here is a case with one single target volume and *two* organs at risk. The three structures are configured in a way that for a given target dose both organs at risk unevitably get some dose as well. Sparing one organ at risk at risk might lead to a higher dose to the other, i.e. there are many Pareto optimal plans with different compromises for the two risks. The questions are:

- 1. What are the Pareto optimal plans?
- 2. How sensitive are these plans to variations of the criteria? E.g. where is the point at which one organ cannot be improved without heavily worsen the other?

6.1.2 Material and Methods

The clinical case used for this study was the head and neck case introduced in chapter 3 and 5. To reduce the complexity of the problem, the brainstem and the boost were taken out of the optimization. Those structures have turned out to be uncritical and not dose limiting. The remaining structures are the target volume (now including the volume of the former boost), the parotid gland and the spinal cord.

In this first incarnation, a simple "brute force" strategy for generating the database of Pareto optimal plans is implemented. Following settings are used:

- The dose prescription for the target volume is set to minimum dose = maximum dose = 60 Gy.
- For the organs at risk, EUD constraints are used. The EUD maximum constraint is varied from 0 to 30 Gy in steps of 2 Gy for both parotid gland and spinal cord. All combinations of constraints are fully optimized, resulting in 16x16=256 plans.
- Each plan is normalized to a target $\text{EUD}(a_{-} = -8)$ of 60.0 Gy. To ensure a certain homogeneity of the target dose distribution, a plan is rejected if the target EUD belonging to $a_{+} = 10$ is higher than 67.0 Gy.
• If the plan is accepted, it is added to the database. The EUD values, the intensity matrices, and an optimization summary is saved to the harddisk.

After normalizing the plans and filtering out the plans with too inhomogeneous target dose, the dose to the target in all remaining plans can be considered as equivalent. This allows the graphical representation of the plan database by plotting the EUD values of parotid gland and spinal cord on a two-dimensional plane.

6.1.3 Results and Discussion

The total calculation time was 17 hrs for the 256 plans (approx. 4 min per plan) on a Pentium IV 1.8 GHz workstation. 88 of these plans were rejected because of the homogeneity criterion. The remaining 168 plans can be divided into the subgroups "Pareto optimal" and "dominated". A Pareto optimal plan can easily be identified by the fact that there exists no other plan with a lower EUD in one organ at risk without a higher EUD in the other organ at risk. It turned out that the database contains 19 Pareto optimal and (168-19)=149 dominated plans.

Figure 6.1 shows the database of plans graphically by plotting the EUD of the parotid gland in dependency on the EUD of the spinal cord. The Pareto optimal plans mark the border between the physically feasible area to the non-feasible area. The planner is only interested in the Pareto optimal plans, the dominated plans are shown only for informational purposes. A final database contains only the Pareto optimal subgroup.

The graphical representation gives detailed information about the concrete planning situation very quickly. E.g., at first sight the planner realize that it is physically impossible to achieve an EUD of 5 Gy in both spinal cord and parotid gland (unless the target dose is scaled below 60 Gy EUD). It also becomes clear at first sight up to which point the EUD in one organ at risk can be improved without significantly worsening the dose in the other. In this case, most probably a treatment planner would decide between the plans marked with A, B and C. The DVHs of these plans are shown in Fig. 6.2.

The decision which plan will finally be used for treatment is up to the planner and can be based on clinical considerations. If all of these plans would lead to



Figure 6.1: Database of all plans that satisfy the target dose EUD and homogeneity constraints. Each plan can therefore be represented by its EUD values for parotid gland and spinal cord. Only the Pareto optimal plans are of interest for the treatment planner.



Figure 6.2: DVHs of Plan A, B and C.

inacceptable high risks of radiation damage, the dose to the target volume has to be reduced. By scaling the whole database, all plans will remain Pareto optimal and the planner can be sure to stay at the physical optimum.

In the future, the algorithm might detect the most interesting part of the database automatically (in this case, the part around plans A, B and C) and produce additional Pareto optimal plans to allow finer plan tuning (see also the subsection "Plan interpolation" below).

In this simple example, the target dose was at a fixed level of EUD value and dose homogeneity. However, it would be of clinical interest to consider target type structures with two criteria: First, the dose level in terms of EUD regarding a parameter a_- , and a second criterion measuring the dose homogeneity (e.g. the EUD regarding a high positive parameter a_+). The dose homogeneity can then be included into the multicriteria decision process as an additional parameter just like the EUD in an organ at risk. This way no plan has to be rejected, and in that sense a case with one target volume and one volume at risk as described in the introduction would already have three criteria. In fact, in clinical practice it is a frequent situation that the balance between the target dose homogeneity and the dose to one organ at risk is the most difficult problem.

As soon as there are more than two criteria to consider, the database cannot be represented by a simple two-dimensional plot. In this case the more elaborate search tool from ITWM Kaiserslautern presented in chapter 2 has to be used. However, to get an overview for the dependency of one structure on another it might still be useful to lock all other criteria and have a look at the twodimensional plot of the remaining two structures.

6.2 Techniques for more complex cases

It is clear that the crude method used for the clinical test case in the previous section is not sufficient for the multicriterial optimization of cases with more organs at risk and probably several target volumes. If all combinations of constraints are fully optimized, the number of necessary optimization runs would increase exponentially, taking by far too much calculation time. Therefore more sophisticated methods have to be implemented to keep the calculation time within clinically acceptable limits. One method could be to restrict the multicriteria optimization to those structures of the treatment plan that are the crucial, dose-limiting elements. Usually, in treatment plans with many structures there are some organs at risk that are easy to keep within tolerance, and these structures can be excluded from the multicriteria optimization process. One example is the brain stem of the head and neck case shown in the previous chapters. Geometric considerations and test optimization runs could be helpful for detecting and excluding those uncritical structures. The complete generation of the Pareto optimal plan database should not take longer than one day, even when, say, six different criteria have to be considered. In addition to the parameter reduction, in the following two techniques are introduced that let this goal appear realistic to achieve. These techniques are not analyzed in full depth, nor are they already integrated into the optimization process. The following should be understood as suggestions for the future work on this project.

6.2.1 Adaptive constraints

As can be seen from Fig. 6.1, most of the plans that are generated by the "brute force" method are not Pareto optimal. The time needed for the generation of a dominated plan is wasted and could be saved if every single optimization run already ensures a Pareto optimal result. A first test suggests that this can be achieved by the technique of adaptive constraints. Adaptive means that the constraints are not constant throughout the whole optimization run. If the algorithm detects that at some point of the optimization a constraint is reached within a very close interval, it lowers the constraints automatically and continues to improve the specific criterion beyond the original constraint. This has to be done with a lowered penalty factor to avoid unwanted violation of one of the other constraints.

Figure 6.3 shows a first result of the adaptive constraints technique. The first plan (non-adaptive) was optimized without adaptive constraints, so the optimization stopped as soon as the prescribed EUD constraints were reached. The second plan (adaptive) was optimized beyond these constraints, so with the same starting EUD constraints a better optimization result in brainstem and spinal cord could be achieved. The dose to the parotid gland could not be improved



Figure 6.3: Adaptive constraints.

beyond the prescription, indicating that the dose to this organ was already as low as possible.

6.2.2 Plan interpolation

A second technique to further reduce calculation time also avoids unnecessary optimization runs. It is motivated by the fact that between two Pareto solutions a third solution can be approximated by linear interpolation. Because in the dose calculation engine the dose depends linearily on the intensities (eq. 3.1), both the dose distribution and the intensity matrices can be interpolated without a new dose calculation. The interpolated plan will lie in the area of feasible solutions. The EUD values have to be recalculated because except for perfect parallel organized organs the EUD does not depend linearily on the dose values. However, the linear interpolated EUD value already gives an upper limit for the exact EUD value. Therefore in Fig. 6.1 all Pareto optimal plans could be connected by lines; these lines give the upper EUD limits of interpolated Pareto solutions.



Figure 6.4: Interpolation of plans.

The question is how near the interpolated solution will be located to the Pareto front, see Fig. 6.4.

This might be answered by adding the technique of adaptive constraints to the interpolation: Starting from the interpolated plan, additional optimization iterations will shift the plan to the Pareto front. One can expect that many iterations can be saved compared to a new optimization run starting from zero.

Plan interpolation can also be useful during the interactive planning session. If the planner cannot decide between two certain plans, plan interpolation can be a tool to produce and present finer resolved plan alternatives on-the-fly.

72 CHAPTER 6. APPLICATION OF MULTICRITERIA OPTIMIZATION

Chapter 7

Summary and Outlook

Multicriteria optimization in inverse radiotherapy planning is a collaborative research project of the German Cancer Research Center (DKFZ), the Fraunhofer Institute for Industrial Mathematics (ITWM) and the Massachusetts General Hospital (MGH). This thesis dealt with some fundamental techniques and questions associated with the multicriterial approach. Because of the high number of optimization runs, the optimization had to be accelerated. This could be achieved through sampling of the dose calculation matrices. The next question was about which criterion should be used to represent each structure in the multidimensional plan description vector. Two models of the equivalent uniform dose were analyzed and compared. It turned out that the new linear max & mean model is a promising alternative to the established generalized mean model. While the optimization with the max & mean EUD is investigated by the collaboration colleagues at ITWM, a new optimization algorithm had to be developed to implement the generalized mean EUD as an optimization constraint. This problem was solved by using the method of projection onto convex sets. Finally, all techniques were utilized to produce a database of Pareto optimal plans of a simple clinical test case. It gives a first impression of the new quality of treatment planning that multicriteria optimization is capable to deliver. Some advanced techniques for more complex cases, namely adaptive constraints and plan interpolation, are briefly introduced as well.

It should be noted that some of the presented material is useful for existing inverse planning systems even outside of multicriteria optimization schemes. The new inverse planning system at DKFZ makes use of the accelerated dose calculation and is faster in all modes of photon treatment optimization. Also the generalized mean EUD constraints can be used for inverse planning in general. The possibility to mix EUD with physical dose constraints is especially useful to smoothly introduce this new concept to existing systems.

This thesis can be considered as a first milestone of the multicriteria optimization project. While fundamental techniques and algorithms could be realized, further work has to be done to fully exploit the possibilities of the multicriteria optimization concept. For example, the techniques presented in chapter 6 have to be included into the optimization engine in order to produce Pareto optimal plan databases for complex cases. Another investigation will have to deal with the question of the planning horizon, i.e. up to which dose and in what resolution plans have to be generated and included into the database. Maybe the integration of TCP/NTCP models will be useful in this context. Upcoming clinical data will further improve the reliability of both the TCP/NTCP and the EUD concept.

In conclusion, the multicriteria optimization concept provides the basis for powerful tools that assists the radiotherapy planner in achieving the ultimate goal: finding the treatment plan that is the best for the patient.

Bibliography

Abramowitz M. Elementary analytical methods. In: Handbook of mathematical functions with formulas, graphs, and mathematical tables. Stegun A, ed. Washington, DC: US Government Printing Office 1964: 10

Bortfeld T, Schlegel W, Rhein B. Decomposition of pencil beam kernels for fast dose calculations in three-dimensional treatment planning. Med Phys 1993; 20(2): 311-318

Bortfeld T, Stein J, Preiser K. Clinically relevant intensity modulation optimization using physical criteria. In: Starkschall G and Leavitt DD, eds. The use of computers in radiation therapy. Medical Physics Publishing, Madison, WI, USA: ICCR, 1997: 27-30

Bortfeld T. Optimized planning using physical objectives and constraints. Seminars in Radiation Oncology 1999; 9: 20-34

Bortfeld T, Schlegel W. Optimization of beam orientations in radiation therapy: Some theoretical considerations. Phys Med Biol 1993; 38: 291-304

Bortfeld T, Thieke C, Kuefer KH, Trinkaus H. New approaches in intensitymodulated radiotherapy - A new optimization paradigm. In: Kogelnik HD, Lukas P, Sedlmayer F, eds. Progress in Radio-Oncology VII. Monduzzi Editore, Bologna, 2002: 251-258

Brahme A. Which parameters of the dose distribution are best related to the radiation response of tumours and normal tissues? In: IAEA, Radiation dose in radiotherapy from prescription to delivery. Proceedings of the Interregional Seminar for Europe, the Middle East and Africa Organized by the IAEA, Leuven, 1994

Bregman LM. Finding the common point of convex sets by the method of successive projections. Dokl Akad Nauk USSR 1965; 162: 487-490

Burman C, Kutcher GJ, Emami B, Goitein M. Fitting of normal tissue tolerance data to an analytic function. Int J Radiat Oncol Biol Phys 1991; 21(1): 123 -135

Cho PS, Phillips MH. Reduction of computational dimensionality in inverse radiotherapy planning using sparse matrix operations. Phys Med Biol 2001; 46: N117-N125

Cho PS, Lee S., Marks II RJ, Oh S, Sutlief SG, Phillips MH. Optimization of intensity modulated beams with volume constraints using two methods: Cost function minimization and projections onto convex sets. Med Phys 1998; **25**: 435-443

Choi B, Deasy JO. The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. Phys Med Biol 2002; 47: 3579-3589

Dale E, Olsen DR. Specification of the dose to organs at risk in external beam radiotherapy. Acta Oncol 1997; 36(2): 129 - 135

Eisbruch A, Ten Haken RK, Kim HM, et al. Dose, volume, and function relationships in parotid salivary glands following conformal and intensity-modulated irradiation of head and neck cancer. Int J Radiat Oncol Biol Phys 1999; 45: 577-587

Emami B, Lyman J, Brown A, et al. Tolerance of normal tissue to therapeutic irradiation. Int J Radiat Oncol Biol Phys 1991; 21(1): 109 - 122

Gubin LG, Polyak BT, Raik EV. The method of projections for finding the common point of convex sets. USSR Comput Math Phys 1967; 7: 1-24

ICRU Report 50: Prescribing, recording, and reporting photon beam therapy. International Commission on Radiation Units and Measurements, Bethesda, Maryland, USA, 1993

Intensity Modulated Radiation Therapy Collaborative Working Group. Intensitymodulated radiotherapy: current status and issues of interest. Int J Radiat Oncol Biol Phys 2001; 51:880-914 Jackson A, Kutcher GJ, Yorke ED. Probability of radiation induced complications for normal tissues with parallel architecture subject to non-uniform irradiation. Med Phys 1993; 20: 613-625

Kuefer KH, Hamacher HW, Bortfeld T. A multicriteria optimization approach for inverse radiotherapy planning. In: The Use of Computers in Radiation Therapy (ICCR 2000), edited by Schlegel W and Bortfeld T (Springer, Heidelberg, Germany, 2000): 26 - 28

Kutcher GJ, Burman C, Brewster L, Goitin M, Mohan R. Histogram reduction method for calculating complication probabilities for three-dimensional treatment planning evaluations. Int J Radiat Oncol Biol Phys 1991; 21: 137-146

Kwa SLS, Lebesque JV, Theuws JC, et al. Radiation pneumonitis as a function of mean lung dose: An analysis of pooled data of 540 patients. Int J Radiat Oncol Biol Phys 1998a; 42: 1-9

Kwa SLS, Theuws JCM, Wagenaar A, et al. Evaluation of two dose - volume histogram reduction models for the prediction of radiation pneumonitis. Radiother Oncol 1998b; 48(1): 61 - 69

Lahanas M, Baltas D, Zamboglou N. Anatomy-based three-dimensional dose optimization in brachytherapy using multiobjective genetic algorithms. Med Phys 1999; 26(9): 1904-1918

Lee S, Cho PS, Marks II RJ, Oh S. Conformal radiotherapy computation by the method of alternating projections onto convex sets. Phys Med Biol 1997; 42: 1065-1086

Lu XQ and Chin LM. Sampling techniques for the evaluation of treatment plans. Med Phys 1993; 20(1): 151-161

Lyman JT, Wolbarst B. Optimization of radiation therapy, III: A method of assessing complication probabilities from dose-volume histograms. Int J Radiat Biol Phys 1987; 13(1): 103 - 9

Milickovic N, Lahanas M, Papagiannopoulou M, Zamboglou N, Baltas D. Multiobjective anatomy-based dose optimization for HDR-brachytherapy with constraint free dterministic algorithms. Phys Med Biol 2002; 47: 2263-2280 Mohan R. Field Shaping for Three-Dimensional Conformal Radiation Therapy and Multileaf Collimation. Sem Radiat Oncol 1995; 5(2): 86-99

Mohan R, Mageras GS, Baldwin B, et al. Clinically relevant optimization of 3-D conformal treatments. Med Phys 1992; 19(4): 933 - 944

Niemierko A, Goitein M. Random sampling for evaluating treatment plans. Med Phys 1990; 17(5): 753-762

Niemierko A, Urie M, Goitein M. Optimization of 3D radiation therapy with both physical and biological end points and constraints. Int J Radiat Oncol Biol Phys 1992; 23: 99-108

Niemierko A. Reporting and analyzing dose distributions: A concept of equivalent uniform dose. Med Phys 1997; 24(1): 103 - 110

Niemierko A. A generalized concept of equivalent uniform dose. Med Phys 1999; 26(6): 1100

Nill S, Oelfke U, Bortfeld T. A new planning tool for IMRT treatments: Implementation and first application for proton beams. In: Schlegel W, Bortfeld T, eds. The use of computers in radiation therapy. Springer, Heidelberg, Germany: ICCR, 2000: 326-328

Nill S. Development and application of a multi-modality inverse treatment planning system. Ph.D. thesis, Ruprecht-Karls-Universitaet Heidelberg, Germany 2001

Preiser K, Bortfeld T, Hartwig K, Schlegel W, Stein J. A new program for inverse radiotherapy planning. In: Starkschall G and Leavitt DD, eds. The use of computers in radiation therapy. Medical Physics Publishing, Madison, USA: ICCR, 1997: 425-428

Pugachev A, Li JG, Boyer AL, Hancock SL, Le QT, Donaldson SS, Xing L. Role of beam orientation optimization in intensity-modulated radiation therapy. Int J Radiat Oncol Biol Phys 2001 Jun 1;50(2):551-60

Wu Q, Mohan R, Niemierko A, Schmidt-Ullrich R. Optimization of intensitymodulated radiotherapy plans based on the equivalent uniform dose. Int J Radiat Oncol Biol Phys 2002; 52: 224-235 Scholz C, Nill S, Oelfke U, Bortfeld T. Application of an advanced photon dose algorithm for inverse IMRT planning. Med Phys 2002; 29: 1335

Siebers JV, Tong S, Lauterbach M, Wu Q, Mohan R. Acceleration of dose calculations for intensity-modulated radiotherapy. Med Phys 2001; 28(6): 903-910

Söderström S, Brahme A. Which is the most suitable number of photon beam portals in coplanar radiation therapy. Int J Radiat Oncol Biol Phys 1995; 33: 151-159

Sodertrom S, Brahme A. Optimization of the dose delivery in a few field techniques using radiobiological objective functions. Med Phys 1993; 20: 1201-1210

Stein J, Mohan R, Wang X, et al. Number and orientations of beams in intensitymodulated radiation treatments. Med Phys 1997; 24: 149-160

Theuws JC, Kwa SL, Wagenaar AC, et al. Prediction of overall pulmonary function loss in relation to the 3-D dose distribution for patients with breast cancer and malignant lymphoma. Radiother Oncol 1998; 49: 233-243

Thieke C, Nill S, Oelfke U, Bortfeld T. Acceleration of intensity-modulated dose calculation by importance sampling of the calculation matrices. Med Phys 2002a; 29: 676-681

Thieke C, Bortfeld T, Kuefer KH. Characterization of dose distributions through the max and mean dose concept. Acta Oncol 2002b; 41: 158-161

Thieke C, Bortfeld T, Niemierko A, Nill S. From physical dose constraints to equivalent uniform dose constraints in inverse radiotherapy planning. Med Phys 2002c, submitted.

Trott KR. Dose-volume effects. Proc Symp Principles and Practice of 3-D Radiation Treatment Planning. Munich, Germany 1996

Wu Q, Mohan R, Niemierko A, Schmidt-Ullrich R. Optimization of intensitymodulated radiotherapy plans based on the equivalent uniform dose. Int J Radiat Oncol Biol Phys 2002; 52: 224-235

Youla DC, Webb H. Image restoration by the method of convex set projections: Part I-Theory. IEEE Trans Med Imaging 198;2 MI-1: 81-94 Zelefsky MJ, Fuks Z, Happersett L, Lee HJ, Ling CC, Burman CM, Hunt M, Wolfe T, Venkatraman ES, Jackson A, Skwarchuk M, Leibel SA. Clinical experience with intensity modulated radiation therapy (IMRT) in prostate cancer. Radiother Oncol 2000; 55(3): 241-249

Danksagung

Mein Dank gilt an erster Stelle Herrn Professor Dr. Thomas Bortfeld für seine hervorragende Betreuung. Wichtige Anregungen, konstruktive Kritik, und auch Freiräume waren immer da, wo man sie brauchte.

Herrn Professor Dr. Wolfgang Schlegel danke ich sehr für die Aufnahme in die Abteilung für Medizinische Physik am Deutschen Krebsforschungszentrum, für sein Interesse am Fortgang der Arbeit und auch für die Erlaubnis, im Rahmen der Arbeit eine längere Zeit im Ausland verbringen zu dürfen.

Herrn Professor Dr. Josef Bille danke ich für die Übernahme des Zweitgutachtens. Herzlichen Dank an Herrn Dr. Simeon Nill für die Einweisung in die Geheimnisse des inversen Planungssystems KonRad2 und seine unermüdliche Unterstützung bei allen computerbedingten Problemen. Vielen Dank auch an Herrn PD Dr. Uwe Oelfke für die Diskussionen über inverse Planung und seine Unterstützung als Leiter der Arbeitsgruppe "Physikalische Modelle". Auch bei den restlichen Mitarbeitern der Medizinischen Physik am DKFZ - Klaus, Christian, Jan, Hanitra, Luciana, Bernd, Angelika und allen anderen - möchte ich mich für jedwede gewährte Hilfe und die angenehme Arbeitsatmosphäre bedanken.

Bei Herrn PD Dr. Karl-Heinz Küfer und seinen Mitarbeitern vom Fraunhofer-Institut für Techno- und Wirtschaftsmathematik möchte ich mich für die Beratung bei mathematischen Fragestellungen und die sehr gute Zusammenarbeit innerhalb des gemeinsamen Forschungsprojektes herzlich bedanken.

Major parts of this thesis arised from my stay at the Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School, Boston, MA. I would like to express my sincere gratitude to Professor Dr. George Chen, Professor Dr. Andrzej Niemierko, and all other department members for their warm welcome and all their support which made this visit an inspiring experience.

Herrn PD Dr. Dr. Jürgen Debus und seinen Mitarbeitern von der Klinischen Kooperationseinheit Strahlentherapeutische Onkologie am DKFZ danke ich ganz herzlich für die großzügige Unterstützung in der Endphase der Arbeit.

Zu guter Letzt danke ich Simone, die immer für mich da war.