MASTER'S THESIS

Optimization of Beam Orientation in Intensity Modulated Radiation Therapy using a Genetic Algorithm

Lova Carlsson, Lena Krantz

Department of Mathematics CHALMERS UNIVERSITY OF TECHNOLOGY GÖTEBORG UNIVERSITY Göteborg Sweden 2005

Thesis for the Degree of Master of Science

Optimization of Beam Orientation in Intensity Modulated Radiation Therapy using a Genetic Algorithm

Lova Carlsson, Lena Krantz

CHALMERS | GÖTEBORG UNIVERSITY



Department of Mathematics Chalmers University of Technology and Göteborg University SE-41296 Göteborg, Sweden Göteborg, June 2005

Matematiskt centrum Göteborg 2005

Abstract

Every year about 50 000 Swedes are diagnosed with cancer and only about half of them are cured. It is therefore of great importance to improve the cancer treatment. At Sahlgrenska University Hospital Intensity Modulated Radiation Therapy, among other treatments, is used when treating cancer in the head-and-neck region. In this thesis the beam orientations in Intensity Modulated Radiation Therapy are optimized for improved dose plans, with increased tumor control and decreased risk of negative side effects. This is done with a Genetic Algorithm using *a priori* knowledge about the goodness of the beam orientations. The algorithm has been tested on two test cases of head-and-neck cancer, and the results have been compared with a reference plan. The results show that the dose to some sensitive organs can be decreased while the same probability for tumor control is achieved as for the reference plan. The conclusion is that optimizing the beam orientations can give a better dose plan and consequently an improved quality of life for the patient.

Sammanfattning

I Sverige diagnostiseras cirka 50 000 personer med cancer årligen, varav endast drygt hälften botas. Det är därför av stor betydelse att förbättra dagens cancerbehandlingsmetoder. På Sahlgrenska Universitetssjukhuset används bland annat intensitetsmodulerad strålbehandling vid behandling av cancer i huvud-halsregionen. Det här examensarbetet går ut på att optimera strålriktningar vid intensitetsmodulerad strålbehandling för att öka sannolikheten för tumörkontroll samt minska risken för negativa biverkningar. För detta används en genetisk algoritm med *a priori*-kunskap om vinklarnas effektivitet. Tester har gjorts på två testfall av cancer i huvud-halsregionen och resultaten har jämförts med en referensplan. Resultaten visar att en minskning av dosen till vissa känsliga organ kan uppnås utan att försämra tumörkontrollen. Slutsatsen är att en optimering av strålriktningarna kan ge en bättre dosplan och därmed en ökad livskvalitet för patienten.

Preface

This project is a master's thesis in applied mathematics at the Department of Mathematical Sciences at Chalmers University of Technology and Göteborg University. It is part of a larger project on the initiative of the Department of Radiation Physics at Sahlgrenska University Hospital in collaboration with Varian Medical Systems Finland Oy. Our part of the project has been to evaluate the effects of optimizing beam orientations for IMRT. For this purpose we have developed a computer program for beam angle optimization that utilizes the IMRT optimization engine in the Eclipse Treatment Planning System developed by Varian Medical Systems. The other part contained an evaluation of our Beam Angle Optimization program done by Fredrik Nordström, Master's student of Medical Physics at the Sahlgrenska Academy at Göteborg University. Fredrik has also provided us with relevant test cases as well as developed the final dose plans that we have presented in our results. We call the test cases Patient 2 and Patient 3 to be consistent with Fredriks report, that includes an additional test case, Patient 1.

Acknowledgments

We would like to thank our supervisor Professor Michael Patriksson at the Department of Mathematical Sciences at Chalmers University of Technology and Göteborg University. You always cheer us up with your positive encouragements when everything seems hopeless. You are always there to answer our questions even when we are troublesome and you always explain in a pedagogic way, just like one wishes every supervisor would.

We would also like to thank the Department of Radiation Physics at Sahlgrenska University Hospital for investing in a computer for us to use, and especially our supervisors: Senior Medical Physicist Ph.D. Karl-Axel Johansson and Ph.D. Anna Bäck. Karl-Axel for your expertise, quick response and straight answers to our questions. Anna for being critical and kind, and for believing in us. Both of you are always helpful and accommodating.

We would also like to thank Varian Medical Systems Finland Oy, for providing the programming interface with access to the IMRT optimization module, and Research Scientist D.Sc. Katja Pesola for providing us with information about the Eclipse Treatment Planning System and advice on Beam Angle Optimization from your experiences, and for technical support.

We would also like to thank Fredrik Nordström for all the help developing the final dose plans and coping with us and all our wishes.

Finally we would like to thank Eric and Chris for your love and support. Thank you for taking care of us when we have been exhausted after working long days and cheering us up when we have been down after encountering problems. "The Quality of Radiation Therapy: Irradiate the right target, with the right dose, within the right time period."

Karl-Axel Johansson

Contents

1	Introduction	1
2	Radiation Therapy2.1Intensity Modulated Radiation Therapy2.2The Multi-Leaf Collimator2.3The IMRT software Helios	1 3 3 5
3	The dose optimization algorithm used in Helios	6
	3.1 The objective function for the dose calculation	7
	3.2 The Conjugate Gradient Method	0
4	The method for Beam Angle Optimization	9
	4.1 The Genetic Algorithm for Beam Angle Optimization	10
	4.2 The Beam Angle Optimization program	11
	4.2.1 Functions defined in the genetic program	12
	4.2.2 Genetic operators	12
	4.3 Parameter settings for the BAO program	14
5	Test cases for the Beam Angle Optimization	15
6	Results	16
	6.1 Results for Patient 2	16
	6.2 Results for Patient 3	19
7	Discussion	22
A	Dictionary	27
B	Abbreviations	28
С	Descriptive figure of the genetic operators	29
D	Simplified dose calculation algorithm	30

1 Introduction

Cancer is a common disease; just in Sweden about 50 000 are diagnosed with cancer every year, and only about half of them are cured [1]. That is why improving the cancer treatment is important. One improvement to be made is to optimize beam orientations in Intensity Modulated Radiation Therapy, IMRT, and that is the subject of this thesis.

Cancer is a malignant tumor that is caused by unrestrained cell reproduction. Radiation therapy, together with surgery and chemotherapy, are the most important methods for cancer treatment today, and often a combination of the different methods is used. The goal of radiation therapy is to obtain the prescribed dose to tumor cells, while sensitive organs receive as low dose as possible, that is, to kill cancer cells without negative side effects. To achieve the goal of radiation therapy, a treatment called Conformal Radiation Therapy, CRT, is used. This therapy creates a high dose volume, that closely conforms to the shape of the target volume in 3 dimensions, while also minimizing the dose to normal structures [2]. One efficient conformal therapy method is IMRT first introduced by Brahme, 1988 [2]. The benefits of IMRT, compared to traditional CRT, are: less manual work in the development of the treatment plan and more well defined, high dose volumes that conforms to the target volumes with low doses outside. These benefits come from the iterative optimization of the plan being computerized and the use of intensity modulated radiation beams.

This thesis is made on behalf of the Department of Radiation Physics at Sahlgrenska University Hospital, Sweden. At Sahlgrenska University Hospital the software Helios (Varian Medical Systems Finland Oy) is used for the IMRT when treating cancer, mainly in the head-and-neck region. Many IMRT treatment planning software, including Helios, lack algorithms for optimizing the number of radiation fields and their orientation and directions. This research area is highly topical, and there are several recent articles with different approaches written on this subject. The aim of this study is to evaluate the effects of optimizing the beam orientations in IMRT by developing a computer program for Beam Angle Optimization, BAO, and incorporating it in the existing software, Helios. Our approach is as follows: with a Genetic Algorithm, incorporating prior knowledge about the qualities of beam orientations, using Helios for the dose calculations, find the optimal beam configuration for IMRT.

In the next section some background about radiation therapy, including some different techniques and methods that are used today are presented. Then, in Section 3, the algorithm used in Helios is presented and its advantages and disadvantages are discussed. In Section 4 our Genetic Algorithm and the computer program for BAO are presented. Finally, in Sections 5–7 the tests and results achieved with the BAO program are presented and discussed.

2 Radiation Therapy

Radiation therapy is the primary treatment method for treating cancer in the head-and-neck region. This region is difficult in a surgery technical aspect, due to cosmetic issues and to the large amount of sensitive organs. Radiation therapy is based on radiation causing irreparable damage to tumor cells, resulting in cell death. The basis for the success of this type of treatment is that cancer cells are

more sensitive to radiation than normal, healthy cells that recover faster.

Before treating with radiation therapy a Computed Tomography, CT, scan is made to get information about the locations of tumors and sensitive organs. On the CT slices Gross Tumor Volumes, GTV's, i.e. visible tumors, and Organs at Risk, OR's, as well as a larger volume around the tumor called the Planning Target Volume, PTV, are delineated by the oncologist. The PTV is the GTV and regions with high probability for tumor cells, plus added margins for set-up errors, patient motion, linear accelerator alignment errors and other uncertainties. The oncologist also states dose tolerances for the OR's and normal tissue and prescribes dose to the PTV's, partially based on experience but also physical parameters such as the type of tissue, the tumor proliferation, cell density and repopulation [3]. One PTV can be prescribed two or more different doses, for example a high dose to the GTV and a lower dose to the rest of the PTV.

When treating with radiation therapy the irradiation is divided into a number of fields, usually between five and nine. The fields are positioned at a fixed distance around a center point called the isocenter and delivered by the radiation source moving stepwise around the patient, see Figure 1. The isocenter is defined by the planner and is typically located in the center of the main tumor volume. The treatment is divided into fractions with equal dose, and the dose-fractions are delivered daily until the desired dose is obtained, giving the healthy cells time to recover from the absorbed dose while cancer cells receive irreparable damage. The prescribed dose for a PTV, in the head-and-neck region, is usually around 65–75 Gy, where 1 Gy equals 1 J/kg absorbed dose. Most commonly the patient is treated with 1–2 Gy per fraction, daily for four to seven weeks. Every fraction takes 15 to 20 minutes, where each field takes 1-2 minutes to deliver.



Figure 1: The fields are delivered one by one and between each irradiation the gantry rotates to the next predefined angle.

There are basically two kinds of OR's, with serial and parallel structure, to consider when treating with Radiation Therapy. If a serial organ, such as the spinal cord or the optic nerve, receives too high a dose at any point, then the function of the organ is impaired. Parallel organs have a spare capacity, so that a high dose to one point will not affect its functions; but overdosing a larger fraction of the organ will. Organs with mainly parallel structure include the parotid gland, the inner and middle ear, the eye and the lung.

Today, the most common radiation therapy technique is external beam irradiation, using a linear electron accelerator. The most commonly used external beam treatment planning method is forward treatment planning, an iterative process based on trial and error, used for example in traditional Conformal Radiation Therapy. The planner starts by defining the number of beams, their orientations, shapes and static beam intensities and thereafter calculates the dose distribution. Changes in the number of beams, their directions and intensities are made until a desired plan is achieved.

A newer, more efficient method is inverse treatment planning, where the user instead defines the desired dose distribution together with the number of beams and their orientations, then the program calculates beam intensities and shapes that best satisfy the desired dose distribution. The inverse treatment planning technique focused on in this thesis is IMRT.

2.1 Intensity Modulated Radiation Therapy

IMRT is an advanced form of 3-dimensional Conformal Radiation Therapy, that is especially suitable when treating complex cases of cancer, for example with concave PTV's [2]. IMRT utilizes variable beam intensities that can be achieved by a Multi-Leaf Collimator, MLC, compared to traditional CRT that uses uniform beams. The intensity distributions are most often determined by computer optimization techniques for IMRT, compared to manually for CRT. The IMRT treatment plan is both more conformal and with a higher target dose homogeneity than a traditional CRT treatment plan [2].

Optimized IMRT treatment plans have steep dose gradients and are therefore more sensitive to patient set-up errors than conventional treatments [2]. It is therefore extra important that the patient lies in the exact same position every time. To facilitate this, different patient set-up devices, such as, vacuum cradles, plaster casts, face masks or a stereotactic body frame are tried out before treatment.

The desired dose distribution is given to the optimization algorithm in terms of dose constraints and weight factors. The objective is to minimize the total weighted difference between the calculated and the prescribed dose in all points, voxels. The choice of prescription doses and tissue weighting factors are subjective decisions on the part of the treatment planner.

If any of the constraints are violated after the optimization, the treatment planner has to consult the oncologist to revaluate the dose constraints and their weights. In complex cases, the oncologist might have to choose between full tumor control and sparing an OR, when both can not be achieved simultaneously. After the optimized solution is approved by the user, a leaf motion file is produced. This file considers the physical restrictions for the MLC and generates the practically viable solution that comes closest to producing the requested dose distribution.

2.2 The Multi-Leaf Collimator

To achieve the modulated intensities in IMRT a Multi-Leaf Collimator is used. It is placed perpendicular to the rays and acts like a screen. It consists of several adjoining leaves on both sides of the collimator that, depending on their position, will screen off certain rays in the field. Each side of the collimator also has a jaw that will set the field width, see Figure 2.

A MLC can be used in the step-and-shoot technique where the dose from several subfields are added. This is done by first setting the leaf positions and the beam intensities, irradiating and then turning off the beams for rearrangement.



Figure 2: A descriptive figure of a Multi-Leaf Collimator. The static jaws set the field width, the dynamic leaves move across the field, forming a gap where the radiation passes through.

The number of subfields can vary from a few to over one hundred for more complex cases. A less time consuming way to irradiate, is using a MLC in dynamic mode as for IMRT treatment at Sahlgrenska.

A dynamic MLC, dMLC, intensity pattern is a series of many 1-dimensional strips of intensity profiles, where each strip is delivered by one pair of leaves [2]. These intensity strips are achieved by the sliding window technique, which is a dMLC technique in which the gap formed between each opposing pair of leaves varies when traversing across the tumor volume while irradiating.

One undesirable effect associated with the MLC is the tongue and groove effect, see Figure 3. The leaves are designed with overlaps, one side having an extended portion, tongue, and the connecting side of the adjacent leaf having an indented portion, groove. This design is to reduce radiation leakage between two leaves but it contributes to a variable transmission through the tongue, center or groove part of the leaf. The effect is negligible when two adjacent leaves are sideby-side, but produces an under- or overdosed region near the edge of the leaf, when they are apart. Variable beam transmission through rounded leaf ends or through tongues and grooves should be accounted for in IMRT. Radiation leakage through MLC leaves can amount for several percent of the total dose.

The MLC manufactured by Varian Medical systems (Palo Alto, CA) is constrained by a maximum leaf speed of $2.5 \ cm/s$ and maximum field width, without carriage movement, of approximately 14.5 cm [2]. If the field is wider than the maximum field width, the irradiation of the field is done in two or more carriage steps.



Figure 3: A descriptive figure of the tongue and groove effect with a cross-section of the leaves. The groove and tongue profiles are produced by the white and the striped leaf blocking the beamlets, respectively, and the combined profile, when the two leaves are side-by-side.

2.3 The IMRT software Helios

Helios is the IMRT software in the treatment planning system Eclipse, developed by Varian Medical Systems, that is used at Sahlgrenska University Hospital. Since 2001 about 40 patients have been treated using Helios at Sahlgrenska.

Before optimizing in Helios, each field is divided into beamlets, with a width of $2.5 mm \times 2.5 mm$, and the dose deposition coefficient matrix is calculated. The coefficients for each voxel account for anatomical and geometrical information and are associated with the dose contribution from each beamlet. The beamlets are then grouped into leaf-wide groups, since it is practically impossible to modulate the intensities of the beamlets under the same leaf [4].

During the optimization the beamlet weights are adjusted by Helios, in an iterative process, to minimize the difference between the prescribed and the delivered dose distribution. Some lateral scatter, radiation that is spread to the sides, up to 3 *cm* width from the beamlet's path is included in these calculations. This is where the main contribution of the scatter appears but scatter can reach up to 20 *cm* from the beamlet's path. After the optimization the fluences are converted into a leaf motion file, to find the physically viable solution for the dMLC, and the full scatter dose calculation is performed.

Like many other IMRT software Helios lacks algorithms for optimizing the number of fields, their orientations and directions, that is, their gantry and collimator angles and table rotations [5]. The user must define them based on earlier experience and studies. At Sahlgrenska an equi-spaced configuration of seven or nine coplanar beams, beams lying in the same plane, is used.

It has been shown by Samuelsson and Johansson [6] that the treatment plan improves when the number of fields increases; the greater the number of fields, the steeper dose gradients could be achieved. To achieve an acceptable dose homogeneity five fields or more were required. The disadvantage of increasing the number of fields is that the treatment time per fraction, the time that the patient has to lie still, also increases.

It has also been shown that optimizing the beam orientations can improve the treatment plan, especially for complex cases such as head-and-neck tumors [7, 8, 9, 10, 11]. A treatment plan with optimized beam orientations with few fields can in some cases be equal to, or even better than, a plan without a beam orientation optimization with more fields.

Different IMRT software behave in different ways. Therefore, knowledge of the behavior of the system for different choices of dose constraints and weight factors is important for optimal use of the system. For example, in Helios there is no option of choosing hard constraints, that is, a constraint that can not be violated under any circumstances. Therefore it might be necessary to set a dose constraint even lower than the actual dose limit to an OR, in order to receive a clinically acceptable dose distribution. It has been shown in [6] that changing the weight factor or the dose constraint to an OR has a larger effect than changing the weight or the prescribed dose to a PTV.

When making a dose plan with Helios the planner sees graphically in a Dose-Volume Histogram, DVH, throughout the optimization, how well different constraints are fulfilled for the PTV's and OR's. The planner actively changes constraint doses and weights to adjust the plan until a desired result is reached.

3 The dose optimization algorithm used in Helios

The objective is to find a physically realizable solution that achieves full PTV coverage and dose homogeneity while sparing OR's. Constraints for dose distributions in target volumes and OR's are formulated, and the objective is to minimize the violation of these constraints. Therefore there is a penalty associated with the violation of each constraint and the penalties are weighted relative to their importance.

The constraints for parallel and serial organs are constructed differently. Serial organs have max-dose constraints, and parallel organs have dose-volume constraints often combined with a max-dose contraint. A max-dose constraint penalizes for every point that is overdosed, whereas dose-volume constraints penalize only when a certain volume of the organ is overdosed. In the latter case, the penalty is only added for the number of overdosed points that exceed that volume.

3.1 The objective function for the dose calculation

The objective function is a quadratic function built up by two outer sums, where the first one is a sum over the constraints for all the PTV's and the second one is a sum over the constraints for all the OR's and normal tissue. The objective function is a pure penalty function and the ideal objective value is zero and is only attained when no constraint is violated.

The constraints for the PTV's are in their turn built up by two sums where the first sum controlls that the difference between the prescribed dose and the actual dose is small for the target points, and the second sum strives for dose homogeneity in the PTV's, where a minimum and a maximum limit have been set for the dose.

The constraints for the OR's and normal tissue are also built up by two sums, the first limiting the violation of the dose-volume contraints and the second of the max-dose contraints. Each term of both of the OR constraint sums and the sum of the PTV homogeneity contraint is either the squared difference between the actual dose and the contraint dose in that voxel if the constraint is violated, or zero if it is fulfilled.

The objective function to minimize is:

$$F(\mathbf{x}) = \sum_{m=1}^{N_{PTV}} \frac{1}{N_m} \left(\sum_{k \in V_m} w_m (\mathbf{a}_k^T \mathbf{x} - p_m)^2 + \sum_{k \in V_m} (w_{m,min} \cdot \max\{0, (p_{m,min} - \mathbf{a}_k^T \mathbf{x})\}^2 + w_{m,max} \cdot \max\{0, (\mathbf{a}_k^T \mathbf{x} - p_{m,max})\}^2) \right)$$

$$+ \sum_{n=1}^{N_{OR}} \left(\frac{1}{N_{n,dv}} \sum_{l \in V_{n,dv}(\mathbf{x})} w_{n,dv} \cdot \max\{0, (\mathbf{a}_l^T \mathbf{x} - p_{n,dv})\}^2 + \frac{1}{N_{n,max}} \sum_{l \in V_{n,max}} w_{n,max} \cdot \max\{0, (\mathbf{a}_l^T \mathbf{x} - p_{n,max})\}^2) \right),$$
(1)

where N_{PTV} is the number of PTV's, N_{OR} the number of OR's, index m and n denotes the particular PTV and OR respectively, N_m is the number of target points in PTV m, $N_{n,dv}$ is the number of points in OR n, with dose-volume constraint, that exceed the volume given in the constraint and $N_{n,max}$ is the number of OR points in OR n, with max-dose constraint. Index k denotes a particular target point and index l a normal tissue or an OR point. V_m is the set of points in PTV m, $V_{n,dv}(\mathbf{x}) = \{l : \text{voxel } l \in \{N_{n,dv} \text{ voxels in OR } n \text{ with lowest dose}\}\}$ and $V_{n,max}$ is the set of points in OR n. p is the prescribed dose or the constraint dose, w is the weight of the constraint, and the subscripts min, max and dv denotes the type of the constraint, a represents the dose deposition coefficient vector and $\mathbf{x} \ge \mathbf{0}$ the intensity vector of the beamlets. The quadratic objective function is non-convex due to the \mathbf{x} dependence of the set $V_{n,dv}$, hence there can exist local minima.

3.2 The Conjugate Gradient Method

There exist several methods for solving the minimization problem, some of which are deterministic, some are stochastic [12]. The main difference between the two methods is that the stochastic can escape from a local minimum while the deterministic gets trapped when reaching any minimum. The method used in Helios is a deterministic method called the Conjugate Gradient Method. This method is faster than the Steepest Descent Method because it uses conjugated directions. For the Steepest Descent Method, one descent direction is always orthogonal to the previous one. Hence, if the minimum is in a long valley, it might take many iterations of zigzagging across the valley to reach it. The Conjugate Gradient Method contrarily uses linearly independent directions, so for quadratic functions each new iteration adds a new dimension to the descent direction. Thence for every new iteration i, F is minimized over a subset \mathbb{R}^i that is one dimension higher than in the previous iteration. Accordingly, in the last iteration v, F is minimized over the whole set \mathbb{R}^v . Hence, the maximum number of iterations is the number of dimensions, which is the number of beamlets.

Compared to the Newton Method, that is more efficient than the Conjugate Gradient Method, the Conjugate Gradient Method requires less memory. Since the descent direction contains all the previous descent directions, only the last direction, the new direction and the current point have to be stored.

THE CONJUGATE GRADIENT ALGORITHM:

- 0. Set a starting point.
- 1. Find a descent direction, based on the negative gradient and the previous descent directions.
- 2. Take a step to the minimum in that direction.
- 3. Check if the relative decrease of the function value, compared to the previous iteration, is less than a small number, ε ; if so then terminate, otherwise start over from step 1.

When optimizing the dose distribution with the Conjugate Gradient Method in Helios, the gradient has to be derived from Eq. (1). Since the physical interpretation of the gradient is a vector in the steepest ascent direction and since it is a minimization problem, the negative gradient is used. The gradient vector is:

$$\nabla F(\mathbf{x}) = \sum_{m=1}^{N_{PTV}} \frac{2}{N_m} \left(\sum_{k \in V_m} w_m (\mathbf{a}_k^T \mathbf{x} - p_m) \mathbf{a}_k + \sum_{k \in V_m} (w_{m,min} \cdot \max\{0, (p_{m,min} - \mathbf{a}_k^T \mathbf{x})\} \mathbf{a}_k + w_{m,max} \cdot \max\{0, (\mathbf{a}_k^T \mathbf{x} - p_{m,max})\} \mathbf{a}_k \right) \right)$$
(2)

$$+ \sum_{n=1}^{N_{OR}} \left(\frac{2}{N_{n,dv}} \sum_{l \in V_{n,dv}(\mathbf{x})} w_{n,dv} \cdot \max\{0, (\mathbf{a}_l^T \mathbf{x} - p_{n,dv})\} \mathbf{a}_l + \frac{2}{N_{n,max}} \sum_{l \in V_{n,max}} w_{n,max} \cdot \max\{0, (\mathbf{a}_l^T \mathbf{x} - p_{n,max})\} \mathbf{a}_l \right).$$

The descent direction, h, generated in step 1 in the the Conjugate Gradient Algorithm is:

$$\mathbf{h}^{(i)} = -\nabla F(\mathbf{x}^{(i)}) + \beta^{(i)} [\nabla F(\mathbf{x}^{(i-1)}), \nabla F(\mathbf{x}^{(i)})] \cdot \mathbf{h}^{(i-1)},$$
(3)

where the fraction of the direction in the previous iteration, $\beta^{(i)}$, to be included in the new direction is:

$$\beta^{(i)} = \frac{\nabla F(\mathbf{x}^{(i)})^T \nabla F(\mathbf{x}^{(i)})}{\nabla F(\mathbf{x}^{(i-1)})^T \nabla F(\mathbf{x}^{(i-1)})}.$$
(4)

When the descent direction has been found a minimum along it is calculated. A step is taken to that minimum, see step 2. The iteration continues until the termination criterion is reached, see step 3.

If a minimum found in an iteration contains negative beamlet weights, which are physically unattainable, those are truncated to zero giving the feasible minimum along the descent direction. The next direction of minimization is the negative gradient in that point, $-\nabla F(\mathbf{x}^{(i)})$, since the conjugacy relation of successive directions is no longer valid when the minimum was not reached. Consequently, the final solution may differ significantly from the theoretical optimum, and there is no guarantee that it is the optimal feasible solution, nor that it will be reached in the number of dimensions iterations.

4 The method for Beam Angle Optimization

When optimizing beam orientations for IMRT it is impossible in terms of time to calculate all possible solutions and pick the best one. Using a stochastic algorithm that is faster than a deterministic algorithm and does not get trapped in local minima but finds a near-optimal solution is a good approach.

Several stochastic methods for BAO have been developed in the last few years. Li *et al* [8] present an effective algorithm that selects beams automatically. This is done through two separate processes that are implemented iteratively. In the first process beam angles are selected with a Genetic Algorithm and in the second process a Conjugate Gradient Method is used for the dose optimization, solving a simplified, dose-based objective function. Pugachev and Xing [9] use a simulated annealing beam orientation optimization algorithm, incorporating prior knowledge about the goodness of the angles and only evaluating some configurations, to speed up the process. The higher the measure of goodness for the angle, the more likely it will be included in the trial configuration, and the more likely the trial configuration will be accepted and evaluated by the simulated annealing algorithm. In this way the time spent on optimizing the beam intensity profiles for bad configurations is reduced.

While the previous two methods use a simplified, quadratic objective function when optimizing beam angles, Djajaputra *et al* [10] use a dose-volume based objective function, with a fast simulated annealing algorithm, to facilitate a direct comparison with the clinically developed plan. They also optimize angles in three planes compared to the previous two that only optimize coplanar angles. Schreibmann *et al* [11] use a completely different approach with multiple objectives, decoupling the optimization from the decision making process, when optimizing the number of beams and their orientations and weights. The algorithm produces a set of efficient solutions that all represent different clinical trade-offs. In this way the planner is not required to specify unknown information such as constraint weights, but can choose a satisfying solution out of the set.

The method in this report uses a Genetic Algorithm, as in Li *et al* [8] and a dose-volume based objective function, as in Djajaputra *et al* [10], incorporating *a priori* knowledge, as in Pugachev and Xing [9]. This algorithm and program will be described in the following sections.

4.1 The Genetic Algorithm for Beam Angle Optimization

The Genetic Algorithm originates from an analogy with natural selection [13]. Through different genetic operations such as initialization, selection, crossover, mutation, immunity [8] and in our algorithm also cloning and diversity, a nearoptimal solution is found. Other terms that also have a connection to genetics are the following: a gene, j, is a variable, an individual, r, is a solution, a generation, g, is an iteration, and a population, R, is the set of individuals in one generation.

Despite of the Genetic Algorithm being rather unscientific in its foundation, it is a fairly effective method of scanning through a large number of solutions, where each generation is more fit than the previous. It also gives the advantage of having an entire population, including several good solutions, to choose from after the last iteration.

Using a priori knowledge when choosing the angles for the initial individuals will give us a superior first generation, compared to randomizing, and consequently a faster Genetic Algorithm. We have chosen to use a Figure of Merit, FoM, as a measure of goodness for these angles. The FoM for an angle, j, is set to be inversely proportional to the objective value, F_j , of the dose optimization and is calculated individually for each angle in the set. The set of possible angles, J, is constituted by a number of equi-spaced angles, and the angular interval formed between them is denoted angle-interval.

The number of genes and the population size are predefined and kept constant throughout the optimization. The individuals of the first generation are initialized with genes from the set of angles. The first individuals are set to be equi-spaced, starting at every angle-interval degrees from zero until every angle in the set is represented. Having the configuration used at Sahlgrenska in the initial population, we assure to get a solution that is at least as good as for the configuration used today. The remaining individuals' genes are randomized with the probability proportional to their FoM.

After each new generation is "born", the objective value of the dose optimization is calculated for each individual. The fitness is set to be inversely proportional to the objective value. This is problematic if the objective value $\rightarrow 0$ since then the fitness $\rightarrow \infty$, but it is highly unlikely because of the complexity of the problem.

The individual with the highest fitness in a generation is the best individual. To make sure that the best individual from one generation is kept, cloning is performed. That individual proceeds unchanged to the next generation. The selection of parents, u, is performed where the probability to be selected is positively correlated with the fitness. An individual with a high fitness can be selected to more than one parent couple while an individual with a low fitness might not be selected at all. For any parent couple a random number of genes are swapped, through crossover, so that the two children have a mix of both parents' genes.

Each gene from the parents is represented in the children unless a mutation is performed. A mutation is defined as one gene being changed for a random gene, and it occurs with a small probability. The intention of mutation is to decrease inbreeding, by bringing new genes to the population.

If two genes are equal or within a small angular interval then they can each produce an equivalent dose distribution as the two together. We assume that if two genes are within an immunity-interval, one of them is excessive and due to immunity swapped for a different angle. Immunity also controls that there are no parallel opposed beams in any individual; if so, one is swapped. For a more diverse population no two identical individuals are allowed; if so, one of them is rotated angle-interval degrees clock-wise, with the operator diversity.

The algorithm is terminated when the improvement of the fitness for the best individual is less than a small number, ε , in a predefined number of iterations or after a maximum run time.

OUR GENETIC BEAM ANGLE OPTIMIZATION ALGORITHM:

- 0 Initialization
 - a Calculate the FoM for each angle in the set.
 - b Initialize initial individuals.
 - c Check for immunity and diversity.
 - d Calculate the fitness.
- 1 Clone the individual with the highest fitness.
- 2 Select parents, perform a crossover and possibly mutations.
- 3 Check for immunity and diversity.
- 4 Calculate the fitness.
- 5 If the improvement is less than ε in a predefined number of iterations then terminate; otherwise start over from step 1.

After the BAO is done, the dose distribution for the optimal beam configuration is adequately optimized in Helios, to get the final dose plan.

4.2 The Beam Angle Optimization program

The program for BAO is written in C++. It is built up by a main program, bao.cpp, a parameter file, bao_parameters.h, and a function file for the genetic operators, genetic_operators.cpp. A flow chart of the BAO program is presented in Figure 4, and a descriptive figure of the genetic operations is presented in Appendix C. The bao.cpp program takes the input parameters:

-case	The directory for the Patient files.
-in_individual	A text file with an initial configuration containing no_genes fields (default is a FoM-randomized configuration).
-tot_no_genes	The total number of angles in the set (default is 72).

-no_genes	The number of fields in a configuration (default is 7)
-im_interval	The length of the immunity interval in degrees (default is 10).
-max_run_time	The maximum run time for BAO (default is 900 minutes).
-no_improvement	The maximum number of iterations with no improvement (default is 10).

The program starts by calculating FoM for each angle in the set with the function fom. Then it initializes the first generation with a predefined number of genes for a predefined number of individuals with the operator initialize. immunity is checked for each individual and diversity for the whole population. Then the fitness for each individual is calculated with the function fitness.

A loop for generating new generations starts by randomizing which individuals in the present generation will be parents to the next, through selection. When two different parents have been selected a crossover is performed and possibly a mutation on the children before immunity is checked. cloning is performed on the best_individual, the individual with the highest fitness, and the clone together with the children from crossover are proceeded as the new population to the next generation, if passed diversity. The fitness is then calculated for each individual in the new generation. The program terminates when the improvement of the fitness for best_individual is less than ε , in no_improvement generations, or the max_run_time is reached.

4.2.1 Functions defined in the genetic program

The function fom calculates the FoM for each field in the set with the Varian program, beamlet_optimization.exe. The calculation is done by first writing the field angle to a field parameter file and then executing the program with the command _spawn and the input arguments: the directory with the patient data, the machine data directory, the file with the field parameters, the number of iterations and the file to write the objective value to. FoM for a field is set to be 100 divided by the objective value. The function fom_random_angle randomizes a gene from the set of angles where the probability is proportional to the FoM.

The function fitness calculates the fitness for each individual by first executing beamlet_optimization.exe as above, with the parameter file containing all the genes of the individual, and then setting the fitness to be 1 divided by the objective value.

4.2.2 Genetic operators

The operator initialize generates the first generation of equi-spaced and FoM-randomized individuals. If the user has defined an in_individual it replaces one of the FoM-randomized individuals. The best individual and its fitness proceeds to individual zero in the next generation with the operator cloning. selection randomizes individuals to be parents to the next generation, see Figure 5.

crossover takes two parents and produces two children, by swapping their genes behind a randomized cut. mutation is performed with a small probabil-



Figure 4: A flow chart of the BAO program.



Figure 5: When selecting parents to perform crossover, a number between 0 and the total squared fitness is randomized. The individuals are lined up in order of their index, and each individual is represented by a section of the line equal to its squared fitness. The individual to be selected as a parent is the one containing the randomized number in its section, in this case individual(3).

ity, swapping a randomized gene in the individual for a new FoM-randomized gene. immunity compares all the genes one by one for each individual, and if it finds two that are too close or opposite one another, it replaces the one with the lowest FoM for a new FoM-randomized gene. The new gene is also checked for immunity. diversity compares all the individuals one by one, and if it finds two that are identical, it rotates one of them angle-interval degrees. The new, rotated individual is also checked for diversity.

4.3 Parameter settings for the BAO program

To get our algorithm as effective as possible we have run several tests changing one parameter at the time to find the optimal settings for our program. We have chosen to restrict the number of angles to every five degrees, limiting the search space to 72 angles, giving $8.5 \cdot 10^{10}$ possible solutions, for configurations of nine genes. We have also found good results when prohibiting adjacent fields and setting the immunity interval to 10 degrees; the genes are hence at least 15 degrees apart. This limits the search space further, lowering it more than tenfold to $6.0 \cdot 10^9$ [14].

We have chosen to run 6 IMRT iterations with Varian's beamlet optimization program, when calculating the fitness for each individual. This is sufficient because the most changes in the objective value occur in the first few iterations, and seem to level out afterwards. When calculating the FoM for each gene in the set we only run Varian's beamlet optimization 3 iterations, which is sufficient for one field, where the most improvements takes place in the first three iterations.

We found that having the probability to be selected as a parent being proportional to the square of the fitness gave better results than to the fitness itself. This is due to the fitness range being small and the probability for the best individual in a population to be selected is not much larger than for the worst individual, thus squaring the fitness widens the range.

For the termination criterion we have made the appraisal that the improvement of the best fitness must exceed one percent of the previous best fitness to have a big enough effect on the objective value; hence ε is set to one percent of the best fitness. The program terminates if the improvement of the solution is less than epsilon in 10 generations. We have also set the default time limit to 15 hours, which is the usual time from getting off work in the evening to starting work the next morning, assuming that it will be run overnight.

The number of genes is set by the planner depending on the complexity of the case. We have set the population size according to Li *et al*'s [8] empirically found population size of double the number of genes. Since we want an odd number of individuals to have an even number of parents plus one cloned individual, we have set the population size to double the number of genes plus one.

We have compared the mean fitness of randomized individuals with FoMrandomized individuals, and the results confirm that using *a priori* knowledge gives a superior first generation. Equi-spaced individuals have a relatively high fitness, thus using equi-spaced individuals in the first generation is also a good choice.

5 Test cases for the Beam Angle Optimization

The BAO program was run on two test cases, Patient 2 and Patient 3, both with cancer behind the nose in the upper part of the throat: Nasopharynx cancer. This is a complicated area to treat, hence the number of genes has been set to nine for both of the test cases, where using seven fields do not give a full PTV coverage.

Both of the test patients were prescribed a treatment of a combination of chemotherapy, external and internal Radiation Therapy, i.e. injection of radioactive material to the PTV. Three different PTV volumes were defined: PTV-0.5 is the GTV with a 0.5 cm margin, PTV-1.5 is the GTV with a 1.5 cm margin and PTV-N is a volume containing lymph nodes, where the probability for microscopic spread of tumor cells is relatively high. For the external radiation therapy the prescribed doses to the three PTV's for both Patient 2 and Patient 3 were:

The constraint doses and weights for the OR's were based on common dose limits for organs in the head-and-neck region, see Table 1. Patient 2 has a relatively uncomplicated Nasopharynx cancer with a small tumor, concentrated to the right side without any major overlaps of OR's, see Figure 7. Patient 3 is a more complicated case with a larger tumor. PTV-N overlaps about 50% of each of the parotid glands and PTV-1.5, that is prescribed an even higher dose than PTV-N, overlaps about 20% of the left parotid gland and about 50% of the right one. It is therefore impossible to spare the right parotid gland without loss in PTV coverage. The constraint weight for the right parotid is therefore set to a low priority, since it is more important to achieve the prescribed dose to the PTV's in the overlapping volume.

Since the BAO algorithm is stochastic, the results may vary for each run, and therefore we have run the BAO four times per case to compare run times and objective values. The results have been compared to a reference plan, which is the plan used for IMRT treatment at Sahlgrenska. The reference plan for Patient 2 is: 5° , 40° , 70° , 120° , 160° , 200° , 240° , 295° , 330° . It was configured by the planner since an equi-spaced configuration starting at zero degrees would have beams passing through the shoulders. The reference plan for Patient 3 is the one

Organ	1% of the volume	Maximum	Mean dose	Priority
~	may exceed [Gy]	aose [Gy]	less than [Gy]	
Spinal cord	46	50		High
Brain stem	54	60		High
Optic chiasm	54	50		High
Optic nerves	54	60		High
Eyes			35	High
Temporal lobes	60	65		Medium
Parotid glands			26	Medium
Ears			50	Medium
Pituitary gland	40			Low
TM joints	70	75		Low
Lenses	10			Low
Oral cavity	55	65		Low
Larynx			45	Low
Mandible	70	75		Low

Table 1: Dose limits and priorities for OR's in the head-and-neck region.

of equi-spaced beams starting at zero degrees, that is the standard configuration used at Sahlgrenska.

With the optimal configuration from one of the BAO runs for each patient we have made a final dose plan in Eclipse. The planner has first made sure to get the prescribed dose to the PTV's and fulfilling the maximum dose constraints to serial organs, and then focused on minimizing the dose to the parallel organs: oral cavity and parotid glands with dose-volume constraints.

6 Results

The results of run time, number of generations, fitness and objective value for all four runs are presented in a table for each case. The first run in each table is the one that we have used to make a final dose plan in Eclipse. The results from this plan are presented in a table and illustrated in a DVH. The FoM's are illustrated in a graph for each patient along with the optimal beam configuration marked on a CT-slice, of the cross-section where the tumor is the largest. The results from the BAO runs were obtained with a 3 GHz, single processor PC.

6.1 **Results for Patient 2**

A diagram of the Figures of Merit for all the angles is presented in Figure 6. Comparing the FoM-diagram with the CT-slice in Figure 7 show that the angles with high FoM have a clear path to the PTV's and for the angles with low FoM, the paths to the PTV's are blocked by OR's. For example angles around 300 degrees have a short, clear path to the PTV's while angles around 100 degrees are blocked by the left parotid gland and the Temporo-Mandibular, TM, joint. The angles for the configuration from the first BAO run are shown in the CT-slice in Figure 7. The angles are spread out around the head, mostly located at



Figure 6: A diagram of the FoM's for each angle in the set for Patient 2. Zero degrees is straight from the front with increasing degrees to the left around the patient.

high FoM's, and the majority appearing at FoM-peaks.

The results from the four runs on Patient 2 are presented in Table 2 and show that there are variations in the run time, but the variations of the objective value are small. The average of the objective values from the four BAO runs is 10.1% lower than for the reference plan.

Table 2: The results for Patient 2 from four different runs of the BAO, their average and the fitness and objective value for the reference plan for comparison.

Run	Run time	Generations	Fitness	Objective Value
1	9 h 49 min	12	132.28	0.00756
2	11 h 22 min	14	132.80	0.00753
3	14 h 40 min	19	138.70	0.00721
4	14 h 40 min	19	136.99	0.00730
Average	12 h 38 min	16	135.14	0.00740
Reference	-	-	121.51	0.00823

The results from the final dose plan for the first run on Patient 2 show that the mean dose to the oral cavity was decreased by 3.9 Gy and by 1.3 Gy and 5.8 Gy to the right and left parotid gland, respectively, see Table 3. The DVH for the PTV's, the oral cavity and the parotid glands confirms that the optimized beam configurations give lower doses to the referred OR's than the reference plan while the dose differences to the PTV's are small, see Figure 8. By optimizing the



Figure 7: A CT slice of Patient 2 seen from below, with delineated PTV's and OR's. The cross represents the isocenter and the arrows represent the angles for the optimized beam configuration from the first run.

OR	Plan	Min[Gy]	Max[Gy]	Mean $[Gy]$
Dody	Optimized	0.0	75.9	18.6
Dody	Reference	0.0	75.1	18.6
Oral Cavity	Optimized	0.0	72.7	42.6
OfarCavity	Reference	0.0	70.7	46.5
Right Parotid	Optimized	13.2	65.7	31.8
	Reference	12.6	63.2	33.1
Left Parotid	Optimized	8.5	55.3	25.8
	Reference	13.0	55.7	31.6

Table 3: Minimum, maximum and mean doses for Patient 2 for the body and three different OR's using the optimized beam configuration compared to the reference plan.

beam angles we have achieved a dose plan that gets equivalent PTV coverage as the reference plan, fulfills the dose constraints for the serial OR's and decreases the dose significantly to the oral cavity and the parotid glands.



Figure 8: A Dose-Volume Histogram for Patient 2, comparing the dose to the PTV's, the parotid glands and the oral cavity between the optimized beam configuration and the reference plan.

6.2 **Results for Patient 3**

The FoM's for Patient 3 presented in Figure 9 show that angles in the one hundreds are disadvantageous while there are peaks around 350° and 65° . The range of the FoM's for Patient 3 is relatively small, from about 100 to 500, compared to that for Patient 2 which has a range from about 150 to 1600. This is due to Patient 3 being a more complex case, and there is no angle that is clearly efficient. This can also be seen when comparing Figure 7 and Figure 10 where the PTV's are larger for Patient 3 and overlapping some OR's while for Patient 2, PTV-0.5 and PTV-1.5 are smaller and bounded to the right side.

The angles for the optimal beam configuration of the first run are presented in the CT-slice of Patient 3, see Figure 10. The configuration is close to being equi-spaced with intervals of 25° to 55° between the angles. The optimal angles do not seem to coincide with high FoM's.

The results from the four BAO runs are presented in Table 4; they show that the average of the objective values from the four BAO runs is 13.7% lower than for the reference plan. The run times show that the BAO program takes up to 15 hours and requires therefore more than a working day.

The minimum, maximum and mean doses to the body and some OR's after the final dose plan is developed, are presented in Table 5. Using the optimal beam configuration decreases slightly the dose to the body, oral cavity and the left parotid gland while increasing the dose greatly to the right parotid gland.



Figure 9: A diagram of the FoM's for each angle in the set for Patient 3. Zero degrees is straight from the front with increasing degrees to the left around the patient.

Run	Run time	Generations	Fitness	Objective Value
1	14 h 25 min	21	89.61	0.01116
2	11 h 59 min	17	86.36	0.01158
3	12 h 34 min	18	89.29	0.01120
4	13 h 18 min	19	91.58	0.01092
Average	13 h 4 min	19	89.13	0.01122
Reference	_		76.92	0.01300

Table 4: The results for Patient 3 from four different runs of the BAO, their average and the fitness and objective value for the reference plan for comparison.

Both the reference plan and the optimized plan give too high a dose to the right parotid gland for it to be spared, and therefore the difference between 54.5 Gy for the reference plan and 57.1 Gy for the optimized plan is insignificant. By decreasing the mean dose to the body, oral cavity and left parotid gland the dose plan has been slightly improved.

The DVH for the PTV's, the oral cavity and the parotid glands is presented in Figure 11. It shows that the reference plan and the optimized plan have an equivalent PTV coverage for all three PTV's. By optimizing the beam angles we have achieved a dose plan that gets the same PTV coverage as the reference plan, fulfills the dose constraints for the serial OR's and decreases the dose to the oral cavity and the left parotid gland.



Figure 10: A CT slice of Patient 3 seen from below, with delineated PTV's and OR's. The cross represents the isocenter and the arrows represent the angles for the optimized beam configuration from the first run.

OR	Plan	$\operatorname{Min}\left[Gy\right]$	Max[Gy]	Mean [Gy]
Body	Optimized	0.0	81.1	21.3
Dody	Reference	0.0	77.8	21.6
Oral Cavity	Optimized	24.3	72.7	52.3
OfarCavity	Reference	21.1	73.1	53.7
Right Parotid	Optimized	15.4	76.4	57.1
	Reference	18.4	76.6	54.5
Left Parotid	Optimized	11.8	75.9	41.8
	Reference	11.8	74.9	42.5

Table 5: Minimum, maximum and mean doses for Patient 3 for the body and three different OR's achieved with the optimized beam configuration compared to the reference plan.



Figure 11: A Dose-Volume Histogram for Patient 3, comparing the dose to the PTV's, the parotid glands and the oral cavity between the optimized beam configuration and the reference plan.

7 Discussion

The main disadvantage of our algorithm is that it is slow, with run times between ten and fifteen hours for nine fields. Today the treatment preparation procedure, from the CT scan to the first treatment, takes about one week. With the BAO program running overnight the planning time should not be effected. Some of the manual work could even be reduced, since the beam configuration is near optimal when the dose plan is developed and it is therefore easier to develop a good plan in Eclipse. For a less complicated case where seven fields are sufficient, that is the most common number of fields used for IMRT at Sahlgrenska, the run times will decrease considerably.

Since the program prints the individuals, their fitness and process time after each generation one can evaluate the progress of the solution. If a time limit of three hours would have been set for the runs of the BAO program then the average of the objective values would have been 2.9% and 5.0% higher than the final objective value for Patient 2 and Patient 3 respectively. The average of the objective values would still have been 5.7% and 9.4% lower than for the reference plan for the same patients. These results conclude that improvements can be made with shorter runs as well.

A large part of the run time for our BAO is spent calculating the objective value with Varian's beamlet optimization program and this is where a lot of time could be saved. Our initial approach was to speed up the calculations by creating a simplified dose calculation program according to the algorithm in Appendix D, to be used during the BAO. We were not able to complete the program for the simplified dose calculation because of lack of time and lack of information about the files containing information from the CT scans and the tumors and OR's delineated by the oncologist. On the other hand the simplified dose calculation

might not be adequate to get a good BAO program since it is less precise. The loss in time is the gain in precision. More research is needed to determine if the simplified dose calculation is adequate for determining good beam orientations.

The number of IMRT iterations in the BAO has significance in finding the best solution. In our limited number of tests we found that 6 iterations was enough, but it is possible that the optimal number of IMRT iterations is higher in more complex cases, where the precision is of greater importance. An increase in the number of IMRT iterations increases the run time. With more extensive tests on the termination criteria, the population size, the immunity interval, the total number of angles in the set, the composition of the first generation, and the probability to be selected as a parent, better settings might be found.

An advantage of our algorithm is that we have several beam configurations to choose from after running the BAO. In our thesis we have only evaluated the best solution from each run, but for a more thorough evaluation one could easily, with a small addition to the code, get a list of the best solutions from the entire run. From these solutions one could choose the configuration achieving the most desirable result after developing the dose plans in Helios.

Our algorithm is solely based on minimizing the objective function, which is a good but not complete measure when finding an optimal beam configuration. A beam configuration with a lower objective value might have hot spots or other trouble areas that cannot be accepted while another plan with a higher objective value might be more adequate. This is partly an effect of the inability to set hard constraints in Helios.

Another effect of the inability to set hard constraints in combination with the fact that it is difficult to estimate the values of the constraint doses and weights, is that adjustments of the constraints must be made successively during the development of the dose plan. Thus the final constraint doses and weights might vary between optimal dose plans for different beam configurations for the same patient. During the BAO there is no possibility to adjust constraints, why setting the constraints before the BAO is of big importance so that the objective function for the BAO is as close as possible to the objective function in the final dose optimization.

The development of a dose plan is very complex and it might not be adequate to base the objective function on physical parameters alone. Many researchers today are developing methods of Biologically Optimized Radiation Therapy, based on the biological effects of radiation, but there is still a lot of research to be made.

After viewing the results from the BAO we have seen that when a solution improves from one generation to the next there is often just a small change in some of the genes. A different approach to BAO could be to make small adjustments on already good configurations, for example equi-spaced or based on FoM's.

Based on the results from the two test cases we make the conclusion that the FoM's have a greater importance in less complicated cases where the PTV's are limited and not overlapping OR's, and clear paths to the target can be found. In a more complex case, where there are major overlaps and no open paths, it seems more important to spread out the fields to not concentrate irradiation through just a few OR's, that then would receive way too high a dose. An equispaced configuration is not a bad approximation for a complex case, but one should at least examine which one of the equi-spaced configurations that is the

most efficient. For Patient 3 the most efficient equi-spaced configuration was the one starting at 15° , with an objective value that is 6.2% lower than for the one starting at 0° .

The conclusions are that optimizing the beam orientations give better dose plans with lower dose to some OR's and consequently a decreased risk of complications such as xerostomia. Even though the program is fairly slow, it is worth running due to the gain on the improved quality of life for the patients.

References

- Cancerfonden-Informations- och stödlinjen Vanliga frågor om cancer, Retrieved May 13, 2005 from URL: http://www.cancerfonden.org/ templates/Pressrelease___670.aspx?full=1#ett.
- [2] Ling C, Leibel S, Fuks Z, Humm J, Koutcher J, Hricak H, Amols H, Spirou S, Chui C-S, LoSasso T, Palm Å, Hunt M, Burman C, Zelefsky M, Chong L, Yenice K, Wolden S, Hong L, McCormick B, Rosenzweig K, Alektier K, Yang J, Jackson A, Yorke E, Mageras G, Chang J, Lovelock M, Germain J, Mechalakos J. A Practical Guide to Intensity-Modulated Radiation Therapy. Medical Physics Publishing, Madison, USA, 2003.
- [3] Karl-Axel Johansson (Ph D.), Senior Medichal Physicist at Sahlgrenska University Hospital, oral and e-mail communication.
- [4] Spirou S, Chui C-S. A gradient inverse planning algorithm with dosevolume constraints. Medical Physics 25(3) 321–33, 1998.
- [5] Bäck A. Intensity Modulated Radiation Therapy in the Head and Neck Region using Dynamic Multi-Leaf Collimation. Vasastadens Bokbinderi, Göteborg, Sweden, 2003.
- [6] Samuelsson A, Johansson K-A. Intensity modulated radiotherapy treatment planning for dynamic multileaf collimator delivery: influence of different parameters on dose distributions. Radiotherapy and Oncology 66(1) 19–28, 2003.
- [7] Hou Q, Wang J, Chen Y, Galvin J. Beam orientation optimization for IMRT by a hybrid method of the genetic algorithm and the simulated dynamics. Medical Physics 30(9) 2360–67, 2003.
- [8] Li Y, A, Yao J, Yao D. Automatic beam angle selection in IMRT planning using genetic algorithm. Physics in Medicine and Biology 49(10) 1915–32, 2004.
- [9] Pugachev A, Xing L. Incorporating prior knowledge into beam orientation optimization in IMRT. International Journal of Radiation Oncology Biology Physics 54(5) 1565–74, 2002.
- [10] Djajaputra D, Wu Q, Wu Y, Mohan R. Algorithm and performance of a clinical IMRT beam-angle optimization system. Physics in Medicine and Biology 48(19) 3191–212, 2003.
- [11] Schreibmann E, Lahanas M, Xing L, Baltas D. Multiobjective evolutionary optimization of number of beams, their orientations and weights for *IMRT*. Physics in Medicine and Biology 49(5) 747–70, 2004.
- [12] Webb S. The physics of Conformal Radiotherapy -Advances in technology. Medical Science Series, Sutton, UK, 1997.
- [13] Reeves C, Modern Heuristic Techniques for Combinatorial Problems. McGraw-Hill Book Company Europe, Guildford, UK, 1995.

- [14] Honeyman L, Seating unfriendly costumers A combinatorics problem. 2002, Making Mathematics: Mathematics Research Teacher Handbook, Retrieved February 7, 2005 from URL: http://www2.edc.org/makingmath /handbook/teacher/presentingYourResearch/presentingYourResearch.asp.
- [15] Chuang K-S, Chen T-J, Kuo S-C, Jan M-L, Hwang I-M, Chen S, Lin Y-C, Wu J. Determination of beam intensity in a single step for IMRT inverse planning. Physics in Medicine and Biology 48(3) 293–306, 2003.

A Dictionary

Beamlet	Small beam of radiation.
Collimate	To screen off and make parallel.
Conform	To correspond in form or character.
Coplanar	Lying in the same plane.
Deterministic	Describes an algorithm in which the correct next step
	depends only on the current state.
Gantry angle	The angle of incidence in one plane around the isocenter.
Gray	The unit for absorbed dose, $1 Gy = 1 J/kg$.
Gross	Visible to the naked eye.
Heuristic	Technique which seeks good solutions in a reasonable time.
Isocenter	The center point that the fields are oriented around, usually
	located in the main tumor.
Irradiate	Expose to radiation.
Larynx	The upper part of the respiratory passage containing elastic vocal cords.
Lateral scatter	Spread out to the sides from the median axis.
Mandible	The lower jawbone.
Medulla	The lower or hindmost part of the brain; continuous with the
	spinal cord.
Nasopharynx	The upper part of the throat continuous with the nasal
	passages.
Natural selection	A natural process resulting in the evolution of organisms best
	adapted to the environment.
Oncologist	A physician that practices the branch of medicine that deals
	with cancer tumors.
Optic chiasm	The crossing of the optic nerves from the two eyes at the base
	of the brain.
Parotid glands	Glands situated below and in front of each ear, producing
	saliva.
Pituitary gland	The master gland of the endocrine system; located at the base
	of the brain, also called hypophysis.
Proliferate	To grow or multiply by rapidly producing new cells.
PTV-0.5	GTV plus a 0.5 cm margin.
PTV-1.5	GTV plus a 1.5 cm margin.
PTV-N	PTV containing Lymph Nodes with a relatively high
	probability for microscopic spread of tumor cells.
Stochastic	Statistics; being or having a random variable.
Temporal lobes	The lower part of the brain containing the sensory center of
	hearing.
TM joints	The joints between the head of the lower jawbone and the
	temporal bone.
Voxel	A box-shaped part of a three-dimensional space. Compare
	with pixel in two dimensions.
Xerostomia	Dry mouth resulting from reduced or absent saliva flow.

B Abbreviations

BAO	Beam Angle Optimization
CRT	Conformal Radiation Therapy
CT	Computed Tomography
dMLC	dynamic Multi-Leaf Collimator
DVH	Dose-Volume Histogram
FoM	Figure of Merit
GTV	Gross Tumor Volume
Gy	Gray
IMRT	Intensity Modulated Radiation Therapy
MLC	Multi-Leaf Collimator
OR	Organ at Risk
PTV	Planning Target Volume
PTV-0.5	Planning Target Volume 0.5 cm
PTV-1.5	Planning Target Volume 1.5 cm
PTV-N	Planning Target Volume Node
TM	Temporo-Mandibular



C Descriptive figure of the genetic operators

Figure 12: A descriptive figure of how one generation, g, develops into a new generation, g+1, through different genetic operations. The best individual from generation g is cloned to the next generation. The individuals are selected to be parents with a probability proportional to their squared fitness. Crossover is performed between two parents, where genes are swapped below a randomized cut. Mutation is performed with a small probability, and a new gene is randomized. Immunity controls that no genes are too close or opposite in an individual, if so, one is swapped for a new randomized gene. Diversity controls no individuals are equal in a population, if so, one is rotated angle-interval degrees.

D Simplified dose calculation algorithm

Chuang *et al* [15] have developed an algorithm for determining beam intensities in a single step for IMRT inverse planning. They estimate the beamlet intensities to be proportional to their Figures of Merit's, FoM. The FoM is defined as the ratio between the total dose to the PTV and the total dose to the normal tissue for a specific beamlet. The beamlet intensities, B_{jm} , are for each beamlet, m, in each field, j:

$$B_{jm} = \frac{D_{\text{off}} + \sum_{k=1}^{N_t \in m} w_k d_k \Delta t_k}{\sum_{l=1}^{N_c \in m} w_l d_l \Delta t_l}, \quad m \in j, \ j \in J,$$

where d_i is the dose from beamlet m to voxel i and Δt_i is the intersection distance between the beamlet and the voxel, $D_{\text{off}} \ge 0$ is an offset value to limit the range of the intensity profile, and w_i is the weight of the constraint for voxel i. All B_j 's are normalized so that the average dose to the PTV for all fields are equal, and set to be 100%. The doses d_i are after normalization called f_{ij} , where index i is for the voxel and index j for the beam, and the total dose to voxel i is denoted D_i .

Only the beam weights, q_j , have to be calculated during the dose optimization, since the beams are pre-calculated. In this way they have been able to save computational time. The disadvantage in this approach lies in the inability to adjust the beamlet intensities continuously during the optimization, since they are set *a priori*, and only the weights of the beams are adjusted during the optimization. The objective for a certain beam configuration r is:

$$\begin{split} \min_{\mathbf{q}} F(\mathbf{q}) &= \sum_{k=1}^{N_t} w_k \left(\sum_{j \in r} f_{kj} q_j - p_k \right)^2 + \sum_{l=1}^{N_c} w_l \cdot \max \left\{ 0, \left(\sum_{j \in r} f_{lj} q_j - p_l \right) \right\}^2 \\ \text{subject to} \quad \sum_{\substack{j \in r \\ q_j \ge 0, \quad \forall j \in r.}} q_j = 1, \end{split}$$