An Optimisation Model for Intensity Modulated Radiation Therapy

Matthias Ehrgott
Department of Engineering Science
University of Auckland
New Zealand
m.ehrgott@auckland.ac.nz

Abstract

Intensity modulated radiation therapy (IMRT) can be very effective in the treatment of primary tumours. The challenge in therapy planning is to develop treatment plans that allow effective treatment of the tumour while at the same time limiting the radiation dose to surrounding tissue and organs at risk.

We present a mathematical model of inverse planning for IMRT. The task of such a tool is to identify an optimal treatment plan, which consists of the irradiation directions and intensity profiles to achieve a high dose in the target volume (tumour) while at the same time protecting organs at risk from dangerous high doses.

To optimise intensity profiles and beam directions we consider underdosing the tumour and overdosing any organ at risk as objectives to be minimised simultaneously. Our model allows the incorporation of optimisation of irradiation directions. This yields a multicriteria integer model of the problem. This model has a (possibly very large) set of efficient solutions, from which a representative subset is calculated and stored in a database. An efficient solution is one in which neither the results for the target volume nor for the organs at risk can be improved without deteriorating the results for another volume of interest. We propose an interactive tool that allows the online selection of a plan through the use of the pre-computed database of efficient solutions.

The proposed model and methodology has considerable advantages over existing software packages and we expect it to improve quality of treatment plans while at the same time reducing planning time. We present some results from 2D-prototype software.

1 Introduction

Intensity Modulated Radiation Therapy (IMRT) is the state of the art technology for radiotherapy treatment of cancer patients. Radiation is delivered from a linear accelerator that can move around the patient body, which is immobilised on the “couch” (Figure 1). Modern equipment using multileaf collimators makes very accurate delivery of radiation to the tumour possible while at the same time protecting any organs at risk from radiation. A multileaf collimator uses metal leaves that can move into the radiation field to block out certain areas of radiation across the beam (Figure 2). This technology, however, has made the planning process much more complex than it used to be, and planning systems do not necessarily support the intensity modulation by multileaf collimators.

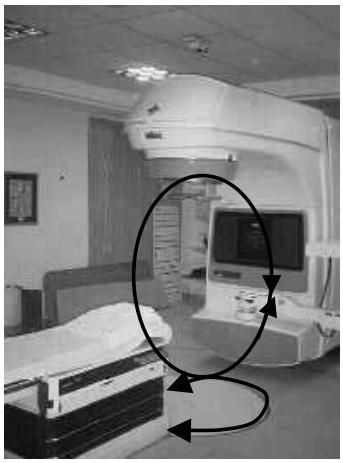

Figure 1: Treatment unit with movements of linac and couch indicated.

Figure 2: Multileaf collimators with movement of leaves indicated.

When choosing a plan for any individual patient, the radiation therapy planner has to determine beam directions, intensity profiles (i.e. the varying intensity across the beam head made possible by multileaf collimation) with the goals of a high (tumouricidal) dose in the tumour, and low dose in normal tissue and organs at risk. In addition, the chosen intensity profiles have to be realised through a number of reconfigurations of the multileaf collimator (the collimator sequence). In the case of step and shoot collimators, these involve irradiating with a given position of leaves, then stopping and reconfiguring the leaves for a new "shot" of radiation. Thus the problem of choosing a sequence of leaf-positions to have a short treatment time arises.

In this paper we focus on optimisation of directions and intensity profiles. We develop a mathematical model for the treatment planning problem. In clinical practice CT or MRI scans of slices of the body along the body axis are used to create a 3D simulation of the body. On these scans, the tumour and organs at risk (called volumes of interest) are outlined and then displayed three-dimensionally in CT simulation programs. At this stage we therefore have K volumes of interest, where an index 1 represents the tumour, and indices $2,...,K$ denote organs at risk. Overlaying a grid on the scans, we then discretise the body into M voxels, $M = M_1+...+M_{K,\text{ each}}$ voxel uniquely identified as belonging either to a specific volume of interest or normal tissue. The smallest possible grid size is given by the resolution of the CT scans and the thickness of the slices. Radiation is delivered from a linear accelerator (linac) that can move around the patient body. The possibilities of these movements plus the possible movements of the couch on which the patient is fixed during treatment define a set of H directions (i.e. possible linac positions relative to the body) from which radiation can be applied. Finally, looking at a beam from a particular direction, the intensity can be modulated across the beam by blocking parts of it out with the collimator leafs. The width of the leaves and the number of positions they can be in determines the resolution of the intensity profile and results in a discretisation of the beam into N bixels per beam direction.

2 A Mathematical Model

For any patient, the oncologist specifies a desired (tumouricidal) dose level L_1 as a lower bound to be achieved in the tumour and tolerable dose levels U_k as upper bounds not to be exceeded in each organ at risk. Finally the number of directions R to be used for irradiation is specified. A small number of directions tends to shorten treatment times, needing fewer movements of the linac, whereas a higher number tends to improve dose distribution and conformance with the target volume.

The mathematical model then uses an intensity vector $x = (x_{11},...,x_{HN})$ of continuous variables, a vector of binary variables $y = (y_1,...,y_H)$ that determine which (at most R) of the H possible directions are used and a deviation vector $T = (T_1,...,T_K)$ that measures deviation of the delivered dose from the desired dose levels in each volume of interest. With a model of beam behaviour that determines the dose absorbed in voxel i given unit intensity in bixel j, the dose distribution vectors $D_k = (D_{k1},...,D_{kMk})$ can be calculated assuming a linear relationship between intensity and absorbed dose. We have used the triple Gaussian pencil beam model of Ulmer and Harder (1995). Here D_{kl} is the dose absorbed by voxel l in volume k. The problem can then be formulated as follows:

$$\min (T_1,...,T_K) \tag{1}$$

$$D_1 = P_1 x \geq (L_1 - T_1)\mathbf{1} \tag{2}$$

$$D_k = P_k x \leq (U_k + T_k)\mathbf{1}, \quad k=2,...,K \tag{3}$$

$$\sum_{h=1}^{H} y_h \leq R \tag{4}$$

$$x_{hi} \leq y_h, \quad i=1,...,N, \; h=1,...,H \tag{5}$$

$$y_h \in \{0,1\}, \quad h=1,...,H \tag{6}$$

$$x_{hi} \geq 0, \quad i=1,...,N, \; h=1,...,H \tag{7}$$

$$T_k \geq 0, \quad k=1,...,K \tag{8}$$

This multiple objective integer programming problem minimises the worst deviation from the desired dose in any voxel of volumes $1,...,K$ (1). (2) represents dose distribution calculation in the tumour and requires that dose to be greater than or equal to the desired dose L_1, minus some deviation T_1. Similarly (3) requires dose in any organ at risk to be less than or equal to U_k plus some deviation T_k. Note that (1) minimises the deviations $T_1,...,T_k$. Here $\mathbf{1}$ denotes a vector of all ones of appropriate dimension. (4) limits the number of beams to be used to R, (5) ensures that only a selected beam direction can be used to irradiate. (6) to (8) define variables to be binary and nonnegative, respectively. The model is based on an earlier one proposed in Hamacher and Küfer 2002.

A solution of multicriteria problem is called efficient, if there is no other solution that is at least as good as the current one in all criteria, and better in at least one. In the IMRT context that means a solution (x,y,T) is efficient, if and only if a decrease in deviation from prescribed dose in the tumour can only be achieved at the expense of an increased deviation for at least one organ at risk, or the deviation for one organ at risk can only be decreased if that for another organ at risk or the tumour can be decreased.

The complete model (1)-(8) is very difficult due to the presence of the binary variables. However, if the irradiation directions are predetermined the binary variables y_h are removed and the model reduces to a multiobjective linear programme. Efficient solutions of an MOLP can be found using multicriteria simplex algorithms. We have used ADBASE (Steuer 2000) for this purpose. Figure 3 shows the (deviation values, i.e. objective function values of) efficient solutions of an example problem, a tumour in the nasal cavity with a prescribed dose of 30 Gy, with two organs at risk: left and right eyeball and optical nerves with a prescribed maximal dose of 10 Gy for both. This example will be the running example throughout the paper, see also Figure 8.

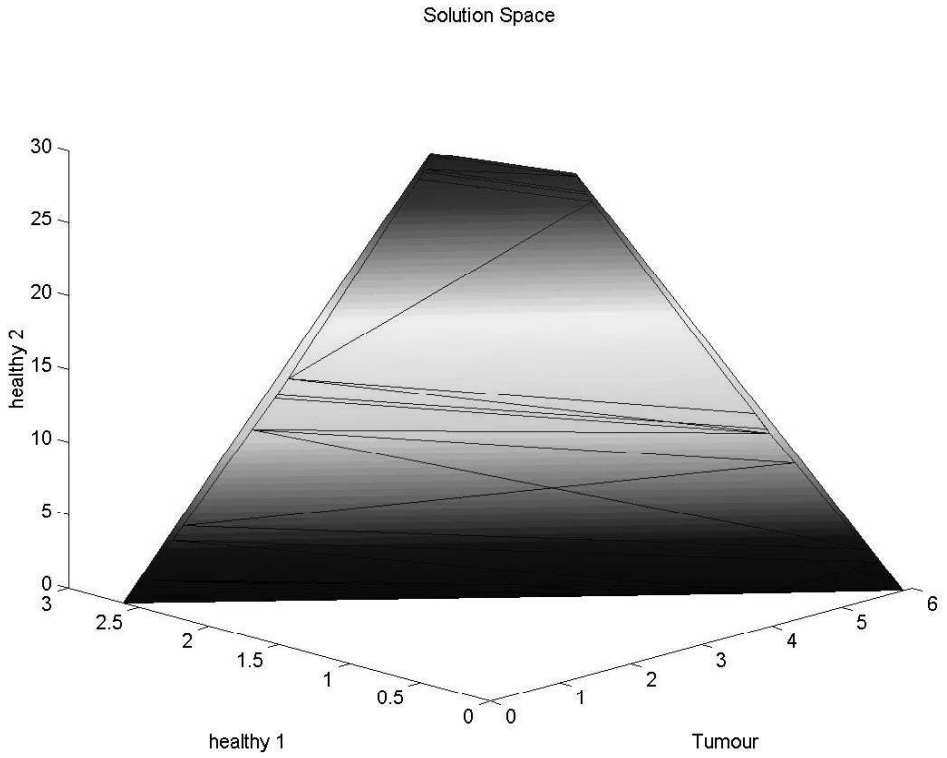

Figure 3: The set of efficient solutions of (1)-(8) for fixed beam directions.

There are infinitely many efficient solutions all representing different treatment plans. However, in practical terms, the deviation values of two points that are close together on the surface shown in Figure 3 are hardly distinguishable. Thus the corresponding treatment plans are hardly distinguishable, too. Therefore, instead of calculating all efficient solutions, or all efficient extreme points that define all efficient solutions, we rather need a representation, that contains distinguishably different efficient solutions that cover the whole efficient set. Such a representation is shown in Figure 4. Note that part of the efficient surface, which corresponded to solutions with a more than 50% overdosing of the right eyeball, has been cut off. Despite being efficient in the mathematical model, such solutions are clinically irrelevant. For the technically minded: this behaviour is due to the underestimation of the Nadir point that is used to evaluate the range of efficient solutions before actually calculating them. An exact computation of the Nadir point is very hard. In this case this inaccuracy produces a positive effect!

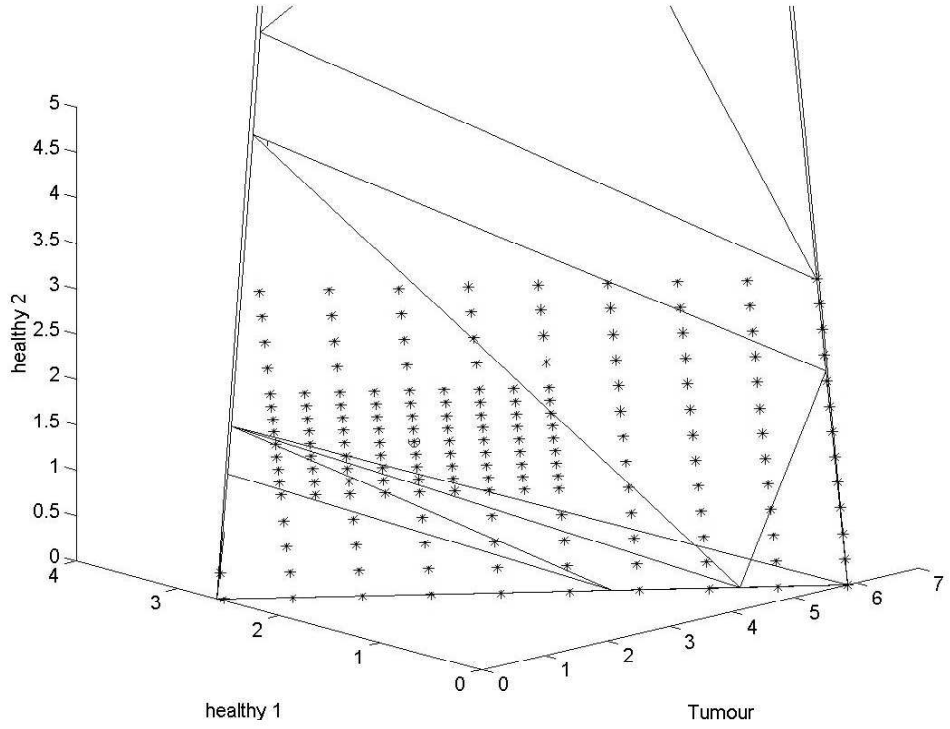

Figure 4: Representative set of efficient solutions for the example of Figure 3.

3 Some Results

In this Section some results are described that we obtained from the model (1)-(8). These concern the use of weights to solve the multicriteria problem and the effect of optimising directions.

For the former, we considered given beam directions and then used equal weights for each objective to get a solution that minimises $1/K(T_1+...+T_K)$. For our standard problem with five equidistant beams the result is shown in Figure 5. This picture clearly shows a typical behaviour of the weighted sum approach: The optimal solution is found at an extreme point, therefore at the boundary of the efficient set in Figure 3. Consequently the results for some volumes (here tumour and right eye) are perfect, and results for others (left eye) are far from ideal. We compared this with a balanced solution that was obtained from minimising *lexmin* $\theta(T)$, where $\theta(T)$ is a nonincreasingly ordered version of the deviation vector T. For details of how this is done see Ehrgott and Burjony (2001). As expected, this solution shows equal deviations for all volumes, and is centrally located in the efficient set. Most existing planning tools use weights of importance.

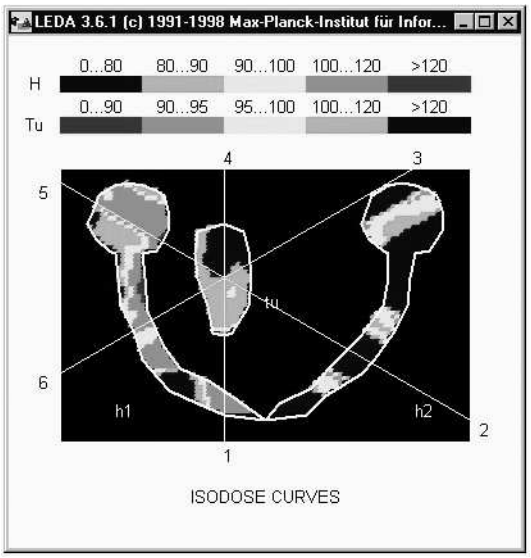

Figure 5: Solution with equal weights, T=(0,1.49,0).

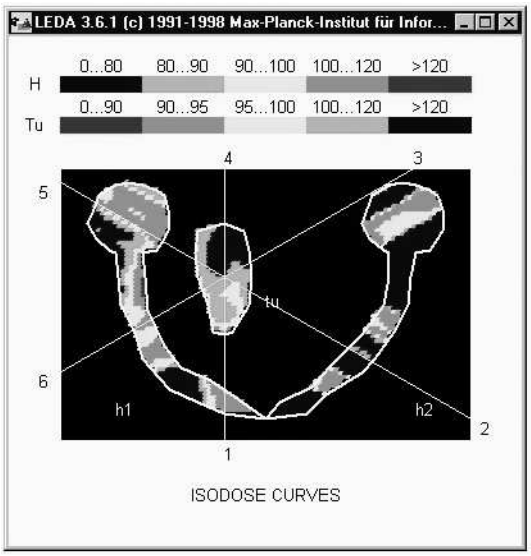

Figure 6: Balanced solution, T=(0.95, 0.95, 0.95).

We studied the effect of optimisation of beam directions using three beams. The effect of beam directions is most significant with few directions, while increasing the number of beams allows for a better spread of radiation over the beams, to reduce damage to organs at risk while still achieving good tumour control. Figure 7 shows the balanced solution for three equidistant beams. Then the scalarised version of (1)-(8) was solved using equal weights for each volume, i.e. minimising $1/K(T_1+...+T_K)$ again, this time including the binary variables for beam direction. Then we solved the balanced problem using the optimal directions found. The solution is shown in Figure 8. The importance of optimisation is obvious: Instead of severely underdosing the tumour as in the left picture, it is now almost completely covered at desired dose level, while results for organs at risk have considerably improved at the same time: we see an 80% reduction in deviation from desired dose foe all three volumes. The picture on the right also shows the counterintuitive features of optimal beam directions: one goes right through the eye, albeit at very low intensity for the central beam, and somewhat higher at the fringe. This beam is needed to cover areas of the tumour that the other two beams cannot. No commercial planning system includes optimisation of beam directions - they all require manual settings.

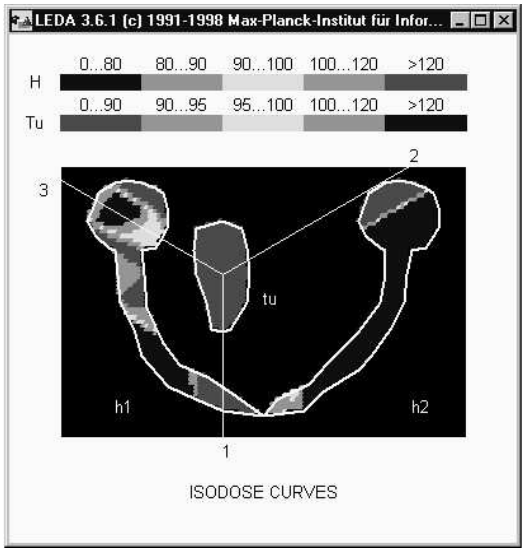

Figure 7: Balanced solution with equidistant beams, T=(5.61, 5.61, 5.61).

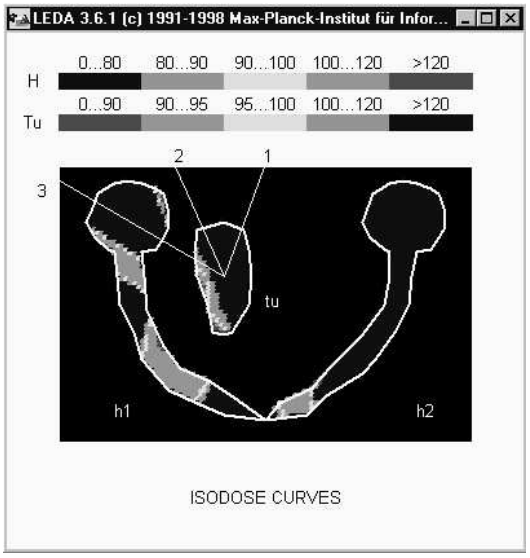

Figure 8: Balanced solution with optimised beam directions, T=(1.72, 1.72, 1.72).

4 Interaction with the User

In this section we want to briefly describe the way the user (the treatment planner or radiotherapist) will work with the optimisation tool. Any radiotherapy planning software needs visualisation tools. CT scans are imported to the planning system, combined into a 3D simulation of the body. Volume contours for the target and organs at risk are outlined and included in the visualisation. The physician then specifies the dose bounds and other relevant parameters for the planning process (these can include dose volume constraints, that more generally than the single dose define which fraction a volume can receive which dose) and inhomogeneity constraints that guarantee a uniform dose distribution in the target volume, which is clinically important.

At this stage the data is passed on to the optimisation module, which computes a database of efficient treatment plans. These correspond to the representative subset illustrated in Figure 4. Once this computation is concluded, the balanced solution, and its treatment plan (intensity profiles) are then exported back to the graphical user interface of the planning tool. If the planner finds the proposed solution unsatisfactory, he can specify changes, e.g. improve result for tumour by 0.5 Gy, but allow 0.2 Gy increase for both organs at risk. A plan that most closely matches these requirements can then be chosen from the database and displayed. This process does not require any further calculations and can be performed online. Because the multicriteria formulation makes trade-offs between results for tumour and organs at risk visible and uses that information, this strategy addresses the issue of assessing the quality of plans with different characteristics that is a problem with many commercial planning systems today. A possible sequence of results could be as follows:

Starting solution:	T=(1.72, 1.72, 1.72)
Improve tumour, keep right eye the same:	T=(1.39, 1.87, 1.72)
Improve tumour, keep right eye the same:	T=(1.07, 2.01, 1.72)
Improve tumour, keep right eye the same:	T=(0.75, 2.15, 1.72)
Improve right eye, keep left eye the same:	T=(0.77, 2.15, 1.58)
Improve right eye, keep left eye the same:	T=(0.79, 2.15, 1.44)
Improve right eye, keep left eye the same:	T=(0.81, 2.15, 1.29)

At this stage a satisfactory result may be obtained. Figure 9 illustrates how the optimisation tool is imbedded in an IMRT planning software system.

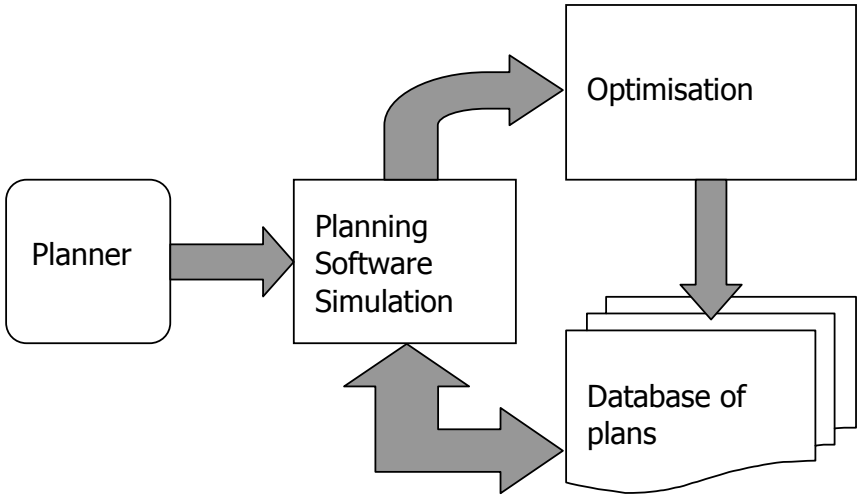

Figure 9: The optimisation tool in a planning system.

In the next step the computed intensity profiles for the three beams needs to be passed on to a collimator sequencer, which determines an optimal sequence of collimator configurations that create the desired intensity profile and that minimises set-up and treatment time. Again, many existing planning systems do use heuristics for collimator sequencing, but do not necessarily find optimal sequences.

5 Conclusion

We have introduced a multicriteria integer programming problem for IMRT treatment planning. Based on this model we have shown how to overcome some of the problems encountered with current optimisation techniques embedded in IMRT planning: the use of priority weights, the lack of beam direction optimisation, and the lack of consideration of trade-offs between overdosing organs at risk and underdosing the tumour.

Acknowledgements

Thanks go to my students, Rick Johnston and Steve McLean, who implemented the direction optimisation and representation of solutions, and to Frank Lenzen and Thang Nguyen, who implemented the original software at the Institute for Techno- and Wirtschaftsmathematik (ITWM), Kaiserslautern, Germany, under supervision of Dr.

K.H. Küfer. Thanks also to Isla Nixon and Lee Dakers of the Physics Section of the Oncology Department at Auckland Hospital for their support and real world feedback.

References

Ulmer,W. and Harder, D. (1995): *A Triple Gaussian Pencil beam Model for Photon beam Treatment Planning*, Zeitschrift für Medizinische Physik, 5:25-30.

Hamacher, H.W. and Küfer, K.H. (2002): *Inverse Radiation Therapy Planning – a multiple objective optimisation approach*, Discrete Applied Mathematics, 118:145-161.

Steuer, R.E. (2000): *ADBASE: A Multiple Objective Linear Programming Solve*, Terry College of Business, University of Georgia, Athens, Georgia, 30602 USA.

Ehrgott, M. and Burjony, M. (2001): *Radiation Therapy Planning by Multicriteria Optimisation*, Proceedings of the 36th Annual ORSNZ Conference, 244-253.