OR Spectrum (2003) 25: 251–264

OR Spectrum

Optimisation of beam directions in intensity modulated radiation therapy planning

Matthias Ehrgott and Rick Johnston

Department of Engineering Science, University of Auckland, Private Bag 92019, Auckland, New Zealand (e-mail: m.ehrgott@auckland.ac.nz; rick@semrits.co.nz)

Abstract. In this paper we consider the problem of selecting optimal beam directions as well as optimal intensity profiles for radiation therapy. Our multiobjective mixed integer programming problem is based on and extends a multiobjective LP formulation for intensity optimisation by Hamacher and Küfer. We use a weighted sum scalarisation to explore the benefits of beam direction optimisation. We propose exact and heuristic methods for solving the problem and present some numerical results.

Key words: Intensity modulated radiotherapy – Multicriteria optimisation – Integer programming – Heuristics

1 Introduction

The application of radiation from external sources (radiation therapy, radiotherapy) is one of the major forms of cancer treatment besides surgery and chemotherapy. The goal of radiation therapy is to deliver a tumouricidal dose to the target volume while at the same time protecting organs at risk from dangerous effects of radiation. In radiation therapy treatment the patient is immobilised on a couch. By movements of the couch and the treatment unit (which can usually rotate 360°) the beam head can be placed in different positions relative to the patient and the target volume. Radiation is then emitted when the beam head is in certain positions around the patient body. In this paper we address the problem of choosing optimal positions of the beam head.

For successful application of radiotherapy it is necessary to achieve dose distributions that conform well to the target volume. The development of intensity modulated radiotherapy (IMRT) and multileaf collimators has improved the success

Correspondence to: M. Ehrgott

rate because they allow shaping the beam and varying intensities across the beam head (intensity profiles). Beam shapes and intensity profiles are realised through multileaf collimation. Multileaf collimators move metal "leaves" into the beam to block out certain parts of the beam. The resolution of the intensity profile depends on the width of the leaves. Various techniques of treatment are possible. Dynamic collimators allow irradiation while the leaves are moving, with step and shoot collimators irradiation is interrupted when leaf-settings are changed. Similarly, the linear accelerators can either allow irradiation from fixed positions or while the accelerator is moving around the patient body (arc treatment). In this paper we do not consider arc treatment.

Traditionally, radiotherapy treatment planning has been done using a trial and error approach: For chosen intensity values and beam directions the dose distribution is calculated and the intensities and/or beam directions are changed, if the dose distribution is not satisfactory. In modern planning systems optimisation models are implemented to achieve good dose distributions. These models use evaluation functions that are usually based on some measure of the deviation from desired dose levels and use weights of importance for the target volume and organs at risk and/or restrictions on dose levels in some of the entities involved. The planning system then uses some mathematical optimisation methods to optimise the evaluation function. Some recent references include [5–10, 14, 15, 17, 20–22, 24].

Because the goals of radiotherapy planning – to achieve a high dose in the tumour while avoiding irradiation of organs and critical structures – are of a contradictory nature Hamacher and Küfer have more recently proposed to consider these goals as separate criteria to be optimised. This idea naturally leads to a multicriteria formulation as given in [4] and also [1] in this issue.

This model is based on a discretisation of the problem. Consider a planning problem with K entities. These entities are defined on CT or MRI scans of the body, which are usually taken in equidistant slices across the area of the body where the tumour is located. Planning systems combine these scans into 3D models of the body on which dose distributions are visualised. We index the target volume with 1, and the organs at risk with $k = 2, \ldots, K$. The relevant body volume is discretised into volume elements, called voxels. The size of the voxels will be given by the distance of CT slices or the resolution of CT images. Let M_k denote the number of voxels in entity k. The dose deposited in voxel i depends on the intensity delivered from a number of beams from R directions. From each of the R positions of the linear accelerator (linac) radiation is emitted according to an intensity profile across the area of the beam head using multileaf collimators. To model these intensity profiles we discretise the beam head into a number of beam elements (bixels). There will be Rn bixels where n is the number of bixels per beam. The size of the bixels will usually be equal to the width of the leaves of the collimator. Assuming that a desired minimal dose L is given for the target volume and desired upper bounds $U_k, k = 2, \ldots, K$ on the dose deposited in the organs at risk the goal is then to find an intensity vector $x \in \mathbb{R}^{Rn}$ such that the dose deposited in the target volume and organs at risk respects the desired bounds.

Due to physical limitations such an x does usually not exist, in particular if R is small. One must therefore find a compromise between underdosing the tumour

and overdosing some of the organs at risk. To formulate an optimisation model it is necessary to have a dose deposition formula. Let p_{ij} be the dose deposited in voxel i at unit intensity in bixel j. The values p_{ij} can be calculated using models of the physical behaviour of radiation, e.g. [19]. Then the dose deposition model becomes $D = Px$, where $D \in \mathbb{R}^M$, $M = \sum_{k=1}^{K} M_k$, is a dose distribution vector. With the objectives of minimising overdosing any organ at risk and minimising underdosing the tumour we can formulate the following multiobjective linear programming model.

$$\begin{aligned} &\min T_1, \ldots, T_K \\ \text{subject to } &D_1 = P_1 x \geq (L - T_1)\mathbf{1} \\ &D_k = P_k x \leq (U_k + T_k)\mathbf{1} \quad k = 2, \ldots, K \\ &x, T \geq 0 \end{aligned} \tag{MOLP}$$

Here, P_k consists of the rows of P corresponding to the voxels of entity k and $T = (T_1, \ldots, T_K) \in \mathbb{R}^K$ is a vector where T_k measures the deviation from the desired dose in the worst affected voxel in each organ, i.e. we measure the deviation by $T_k = ||(P_k x - U_k)_+||_\infty$ for organs at risk and $T_1 = ||(L - P_1 x)_+||_\infty$ for the target volume. $\mathbf{1}$ is a vector of all ones of appropriate dimension. Note that the model can easily incorporate the max and mean model of the EUD concept of Küfer and Thieke [18]. The EUD (equivalent uniform dose) is defined as the uniform dose that has the same effect on an organ at risk as the actual non-uniform dose delivered during treatment. The EUD depends on the organ structure.

Although there is quite some research on the optimisation of intensities, as we have seen above, there is not much research on the optimisation of beam directions. Das et al. [2] present a simple search strategy. Haas et al. [3] try to replicate the approach of a treatment planner using genetic search operators. They perform some trade-off analysis between the quality of dose distributions and the number of beams used. Stein et al. [16] use an exhaustive search and simulated annealing to determine both the number and directions of beams. Pugachev and Xing [12] introduce a scoring function for beam directions and Pugachev et al. [11] use simulated annealing to optimize directions for a given number of beams. Simulated annealing is also used by Rowbottom et al. [13] who compare results with equi-spaced arrangements in a clinical study. A different approach is proposed by Hamacher and Küfer [4] who make use of ideas from location theory. As far as we know there is no commercial IMRT planning system which includes beam direction optimisation. In this paper we present some results on optimisation of beam directions.

The multicriteria model (MOLP) will be the basis of this paper. In Section 2 we extend the model to incorporate the optimisation of beam directions. Then in Section 3 we describe four methods for solving the problem. We present numerical results in Section 4 and conclude with a summary of our achievements in Section 5.

2 Optimisation of beam directions

When we incorporate the beam directions as variables into the intensity optimisation problem the dose deposition matrix becomes dependent on the chosen directions $\theta = (\theta_1, \ldots, \theta_R)$.

$$\begin{aligned} &\min T_1, \ldots, T_K \\ \text{subject to } &D_1 = P_1(\theta)x \geq (L - T_1)\mathbf{1} \\ &D_k = P_k(\theta)x \leq (U_k + T_k)\mathbf{1} \quad k = 2, \ldots, K \\ &x, T \geq 0 \end{aligned} \tag{MONLP}$$

Here the matrix entries $p_{ij}(\theta)$ are nonlinear functions of θ and thus the model becomes a multiobjective nonlinear programming model. We discretise the beam directions. This is justified from the practical point of view, because directions that differ only slightly will cause considerable overlap of beams and therefore lead to very similar plans. In addition, technically it is often the case that treatment units allow only one degree changes of the position of the linac.

In this paper we restrict ourselves to coplanar treatment, i.e. beams are in the same plane as the CT slice. This is the most common treatment technique. Thus, we are working in a two dimensional setting. Note that the model formulation does not change for a full 3D configuration. However, computationally the full 3D problem becomes much larger and harder to solve. We also assume isocentric geometry, i.e. all beams are focused on the centre of the tumour.

With the discretisation of beam directions we can introduce binary variables $y_1, \ldots, y_S$ to model all possible positions of the linear accelerator. Here $S = \lfloor \frac{360}{d} \rfloor$ and d is the angular difference between two consecutive linac positions. Directions which are physically impossible (e.g., irradiating through the couch) can be eliminated to reduce the number of variables. Then let $P_k, k = 1, \ldots, K$ be $M_k \times Sn$ matrices defined as before. Let x_{js} be the radiation intensity of bixel j of beam head s. Then the model can be written as follows.

$$\begin{aligned} &\min T_1, \ldots, T_K \\ \text{subject to } &D_1 = P_1 x \geq (L - T_1)\mathbf{1} \\ &D_k = P_k x \leq (U_k + T_k)\mathbf{1} \quad k = 2, \ldots, K \\ &x_{js} \leq H y_s \quad j = 1, \ldots, n;\ s = 1, \ldots, S \\ &\sum_{s=1}^{S} y_s \leq R \\ &x, T \geq 0 \\ &y \in \{0, 1\}^S \end{aligned} \tag{MOMIP}$$

R is an upper bound on the number of beams to be used in the plan. H is a large number, e.g. the maximum intensity deliverable by the linac. The third constraint guarantees that if an angle is not chosen ($y_s = 0$) then none of the bixels of that beam head position emit radiation.

This problem is a large scale multicriteria mixed integer programming problem, which is very hard to solve. Because the focus of this paper is to explore the benefit of optimisation of beam directions for IMRT planning we do not solve (MOMIP) in the sense of generating all or a representative subset of the Pareto optimal solutions. Here we consider weighted sum scalarisation of objectives into $\sum_{k=1}^{K} \mu_k T_k$ as do current planning systems. Note that for any fixed beam setup (MOMIP) becomes an instance of (MOLP) and all Pareto optimal solutions of that problem can (theoretically) be found using weighted sum scalarisation. However, not all Pareto optimal solutions of (MOMIP) can be discovered using weighted sum scalarisation.

$$
\begin{aligned}
\min \sum_{k=1}^{K} \mu_k T_k & \\
D_1 = P_1 x &\geq (L - T_1)\mathbf{1} \\
D_k = P_k x &\leq (U_k + T_k)\mathbf{1} \quad k = 2, \ldots, K \\
x_{js} &\leq H y_s \quad j = 1, \ldots, n;\ s = 1, \ldots, S \\
\sum_{s=1}^{S} y_s &\leq R \\
x, T &\geq 0 \\
y &\in \{0,1\}^S
\end{aligned}
\tag{MIP}
$$

In principle we could solve this problem by solving an LP for each one of the $\sum_{r=1}^{R} \binom{R}{r}$ sets of directions and choose the best solution. But although in most clinical situations no more than ten beams are used, this approach is too time consuming. In Section 4 we shall see that the problem (empirically) becomes harder as R decreases. It is therefore worth noting that if for some value of R the (MIP) has optimal solution (x^*, y^*) with objective value 0, then we have found an upper bound on the number of beams to be used as well as an optimal solution that achieves all desired dose levels for all $R \geq \sum_{s=1}^{S} y_s^*$, even for all finer discretisations of the angles, i.e. $d' < d$. Note also that usually the objective improves as R increases, see Table 1. Thus, using more directions could still be desirable to allow reduced intensities for each direction.

In the following Section 3 we present a number of approaches for solving (MIP). These are

1. Solving the (MIP) by a commercial solver
2. Local search
3. An LP relaxation heuristic
4. A set covering heuristic.

The first two are integrated methods which attempt to solve (MIP) directly, either optimally or heuristically. On the other hand the latter two are two phase strategies, in which we try to overcome the difficulty of solving the integrated (MIP) model by first determining good beam directions through a direction optimisation model. (Some of) these good directions are then used as input for the intensity optimisation

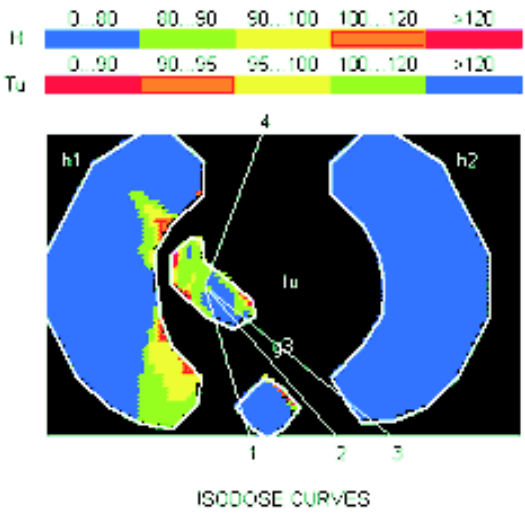

Fig. 1. Optimal (MIP) solution of problem 2

problem. Therefore (MIP) is reduced to an LP, in fact the weighted sum scalarisation of (MOLP). In both methods a sequence of LP's is solved.

3 Solution methods

3.1 Integrated methods

The obvious choice for an integrated method is to simply solve MIP using a commercial MIP solver. We used CPLEX 7.0. It turns out that a feasible solution is often found quickly, but optimality is hardly ever confirmed. Moreover, if R and d are small a feasible solution could not be found after several hours of CPU time. In Table 1 we report optimal objective values and solution times when the number of beam heads is changed for problem 2 (see Section 4 for a description of the test problems). The objective values shown in the table are $\mu_1 \frac{T_1}{L} + \sum_{k=1}^{K} \mu_k T_k / U_k$.

Table 1. (MIP) solution time and objective value versus number of directions R

R	3	4	5	6	7	8	9	10	11	12
Objective value	0.0584	0.0540	0.0500	0.0287	0.0004	0	0	0	0	0
CPU Time	> 500	> 500	> 500	> 500	> 500	134	70.1	13.01	12.01	10.59

Table 1 shows the trade-off between solution quality and number of beams R. Reducing R leads to longer solution times for (MIP) and worse objective values, but usually to shorter treatment times. In Figure 1 we show a dose distribution and directions for the optimal solution of a sample problem that we will refer to

later. Organs at risk are labeled hk, the tumour tu. Dose distributions are given as percentages of desired doses.

In the local search method we first solve LP for a set of R starting directions. We tried equidistant angles $\theta_i = \frac{360}{R} i, i = 0, \ldots, R-1$, random directions with $\theta_i - \theta_{i+1} > 10$ initially, and manual selection. Then each of the directions θ_i is changed in steps of d degrees until θ_{i+1} is reached. In every step LP is solved again. The process is repeated until the objective function does not change anymore.

Algorithm Local Search
$f_0 := \infty$, select θ
Solve LP with θ
$\theta^* := (\theta_1, \ldots, \theta_R), x^* := x, f^* := \sum_{k=1}^{K} \mu_k T_k$
While $f^* < f_0$ do
  For $i = 1, \ldots, R$
    For $l = 1, \ldots, \frac{360}{Rd}$ let $\theta_i := \theta_i + ld$
      Solve LP with θ
      If $\sum_{k=1}^{K} \mu_k T_k < f^*$
        $f_0 := f^*$
        $\theta^* := (\theta_i, \ldots, \theta_R), x^* := x, f^* := \sum_{k=1}^{K} \mu_k T_k$
    Next l
  Next i
End while

We found that different choices of the starting solution affected the solution greatly. The best objective values, however, were very similar for all starting solutions, indicating that the problem has a large number of local minima. Repeated application of the method with random directions never produced the same solution twice.

3.2 Two phase methods

Our first two phase method is based on the LP relaxation of (MIP). In the optimal solution of the LP relaxation some of the y_i will have fractional values. We observe, however, that most of the y_i values are equal to zero. In our tests 10 to 40 y_i variables were nonzero, even when $d = 1$, see Figure 2 for a dose distribution and beam directions for the optimal solution of the LP relaxation of one of our sample problems.

Often, the objective value of the LP relaxation has optimal objective value zero. Therefore it suggests "good" irradiation directions (note that most beams miss critical structures in Figure 2). Notice that the remark concerning an optimal integer solution with objective value 0 on page 255 made above is also valid for an optimal fractional solution. Let $(\tilde{x}, \tilde{y})$ be the optimal solution of the LP relaxation and $\mathcal{R} = \{r : \tilde{y}_r > 0\}$ and $R' = |\mathcal{R}|$. We then compute $w_r = \sum_{j=1}^{n} T_{jr}, r \in \mathcal{R}$ and order the directions $\theta_r, r \in \mathcal{R}$ according to decreasing values of w_r. Next, we order the $\binom{R'}{R}$ subsets $\mathcal{R}_j$ of $\mathcal{R}$ according to $\sum_{r \in \mathcal{R}_j} w_r$ and solve the LP for subsets $\mathcal{R}_j \subset \mathcal{R}$.

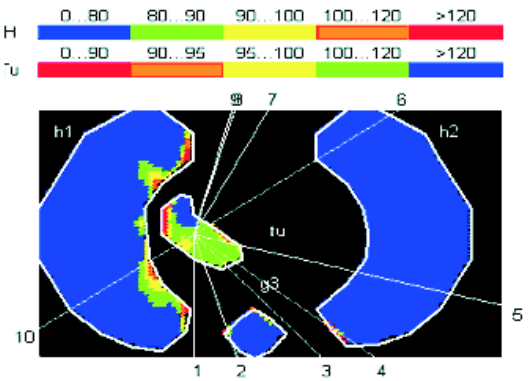

Fig. 2. Optimal solution of the LP relaxation for problem 2

Algorithm LP Relaxation
$f^* = \infty$
For $l = 1, \ldots, \binom{R'}{R}$
  Solve LP with angles $\theta_r, r \in \mathcal{R}_l$
  If $\sum_{k=1}^{K} \mu_k T_k < f^*$
    $\theta^* = (\theta_r : r \in \mathcal{R}_l),\ x^* = x,\ f^* = \sum_{k=1}^{K} \mu_k T_k$
Next l

The algorithm also stops if a preset time limit is reached. If for a time t_0 no improvement of the objective is found, the local search algorithm is called with the current θ as starting solution.

The tests showed that a good solution is found quickly (within five minutes of CPU time), but few improvements were found later, see Section 4 for details. The local search never resulted in directions more than $\pm 20^\circ$ from the original directions indicating that LP relaxation could be used as starting point for steepest descent local search.

The set covering method aims to minimise the damage to the organs at risk while guaranteeing that every voxel in the tumour is irradiated. That leads to the formulation

$$\begin{aligned} \min & \sum_{s=1}^{S} C_s y_s \\ \text{subject to } & Ay \geq \mathbf{1} \\ & y \in \{0,1\}^S. \end{aligned} \qquad \text{(SCP)}$$

Here A is defined by

$$a_{is} = \begin{cases} 1 \text{ if radiation from beam } s \text{ hits target volume voxel } i \\ 0 \text{ else.} \end{cases}$$

It turns out that A is a dense matrix due to the isocentric geometry. The optimal solution of (SCP) always had $\sum_{s=1}^{S} y_s < R$, in fact often a single beam was sufficient to hit every voxel of the target volume. Therefore, we implemented a strategy to solve (SCP) then remove the optimal angles from the problem by setting the corresponding $y_s = 0$ and resolve. This is repeated until a set of promising directions $\mathcal{R}'$ of cardinality $R' > R$ is found. Then we create and solve LPs for subsets of $\mathcal{R}'$.

We tried two ways to generate C_s. C_s is intended to model the damage to organs at risk when using beam s. In the first (weighted angle method) we consider a cone with point at the isocenter and bounded by two half-lines forming $\pm 5^\circ$ angles with the beam direction s. Denote this cone by $cone_s$. Then we considered voxels in organs at risk k that are contained in $cone_s$ for computing C_s.

$$C_s = \sum_{k=2}^{K} \sum_{i=1}^{M_k} \begin{cases} \frac{\mu_k}{U_k} & \text{voxel } i \text{ in } cone_s \\ 0 & \text{else.} \end{cases}$$

In the second (dose deposition) method we consider the actual dose deposition matrices P_k and consider all voxels in all organs at risk.

$$C_s = \sum_{k=2}^{K} \sum_{i=1}^{M_k} \sum_{j=1}^{n} = \frac{\mu_k P_k(i,j)}{U_k}.$$

In Figure 3 we show the C_s values computed by both methods for the problem shown in Figure 2. The directions used in the optimal solution are shown as bold lines (cf. Fig. 1).

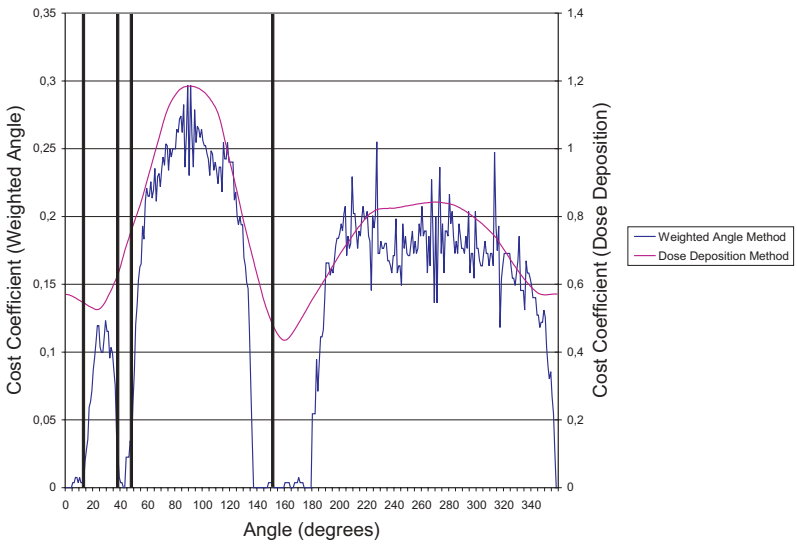

Fig. 3. The objective function coefficients of (SCP) for problem 2. The bold vertical lines indicate the directions used in the optimal solutions

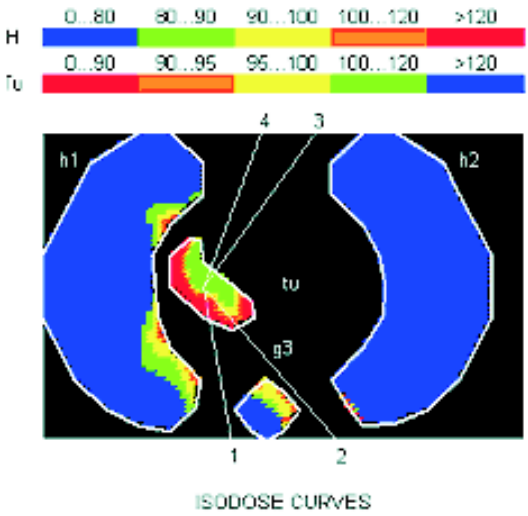

Fig. 4. Best solution found by set covering method for problem 2

The weighted angle method generally makes for a sharper discrimination of "good" and "bad" directions, i.e. low and high values of C_s. The best solution found by the set covering method for the problem of Figure 2 is shown in Figure 4, compare with the optimal (MIP) solution in Figure 1. We note that if d and R' are small the best directions identified by the set covering heuristic will be close together, which may result in bad overall solutions. This observation should be taken into account when the parameters for the method are selected.

4 Numerical results

In Table 2 we summarise results from our experiments. In all tests execution was stopped after 18,000 seconds. Note that the objective values reported here are again $\mu_1 \frac{T_1}{L} + \sum_{k=2}^{K} \mu_k \frac{T_k}{U_k}$. We considered three sample problems. Problem 1 is a tumour in the nasal cavity, with two organs at risk defined by the left and right eyeball and optical nerve, $L = 30, U_k = 10$. Problem 2 is a liver tumour where the organs at risk are left and right lung and spinal chord, $L = 80, U_2 = U_3 = 33, U_4 = 25$. Problem 3 is an artificial one with $K = 6, L = 85, U_2 = 40, U_3 = 45, U_4 = 20, U_5 = 45, U_6 = 20$. In all problems the beams are focused on the center of the tumour (isocentric geometry). In the dose distribution diagrams displayed the tumour is indicated by tu and organs at risk are indicated by hk. The colour coding refers to T_k values as percentages of L respectively U_k.

The results are further illustrated by a number of figures below. The tests were carried out on a Pentium III PC with 512 MB RAM and 900 MHz processor. We used the solver SOPLEX [23] in all methods, except for the MIP results for which we used CPLEX 7.0. The computation times for the MIP and the other methods are therefore not really comparable. With a commercial solver one would expect the computation times for the non-MIP methods to decrease considerably.

Table 2. Numerical results beam direction optimisation

Method	Head 1	Head 2	Head 3	Head 4	Objective	Time
Problem 1, $\mu = (0.3, 0.3, 0.4),\ R = 3,\ n = 10,\ d = 5$						
Set covering	145	185	210	n/a	0.0217	667.40
Local search	155	195	245	n/a	0.0492	3491.18
MIP	165	195	205	n/a	0.0166	604.64
LP relaxation	140	180	200	n/a	0.0269	344.45
Problem 1, $\mu = (0.3, 0.3, 0.4),\ R = 3,\ n = 10,\ d = 2$						
Set covering	96	154	198	n/a	0	16378.50
Local search	160	204	240	n/a	0.0112	7357.01
MIP	64	178	204	n/a	0.0995	3611.84
LP relaxation	138	188	242	n/a	0.0152	6840.51
Problem 2, $\mu = (0.3, 0.3, 0.2, 0.2),\ R = 3,\ n = 10,\ d = 2$						
Set covering	34	50	160	n/a	0.1437	15957.50
Local search	154	184	238	n/a	0.1209	6867.26
Local search	22	162	246	n/a	0.1365	12444.70
MIP	no feasible solution found					
LP relaxation	158	184	242	n/a	0.1210	382.94
LP relax. + LS	158	184	238	n/a	0.1197	400.00
Problem 2, $\mu = (0.3, 0.3, 0.2, 0.2),\ R = 4,\ n = 10,\ d = 2, (d = 1$ for set covering)						
Set covering	0	25	47	147	0.0682	13901.10
Local search	42	156	180	270	0.0455	16980.70
MIP	no feasible solution found					
LP relaxation	156	238	352	10	0.0503	13901.10
Problem 3, $\mu = (0.2, 0.2, 0.2, 0.2, 0.1, 0.1),\ R = 3,\ n = 3,\ d = 2$						
Set covering	42	64	284	n/a	0.0843	11214.00
Local search	132	202	280	n/a	0.0993	3280.71
MIP	no feasible solution found					
LP relaxation	200	244	310	n/a	0.0863	1363.11

Figure 5 shows the improvement of the objective value over time for problem 1. Note that the MIP solver tended to find a good feasible solution but did not find a single feasible solution in some cases. In addition, this first solution was not improved over an acceptable time period. Optimality was hardly ever declared, except when a solution with objective value 0, which is clearly optimal, was found. Figure 6 also shows that all heuristic methods found solutions of comparable quality.

Figure 7 illustrates the importance of beam direction optimisation using problem 1. We compare the best dose distribution obtained with equidistant directions and the optimal solution of the scalarised (MOMIP). We obtained similar results for all three problems considered. The picture clearly shows that the optimisation of directions is very important, especially when R is small. Note that one of the directions in the optimal solution directly intersects one of the organs at risk. This, however, is done at a very low intensity, at the borders of the beam, in order to cover part of the tumour that cannot be reached by the other two directions, whereas the central beam has zero intensity, The low intensity makes it possible to limit the damage to the organ at risk while achieving better tumour control. Thus Figure 7 illustrates that optimal beam directions may be counterintuitive and that the use of optimisation techniques is extremely important to obtain good treatment plans.

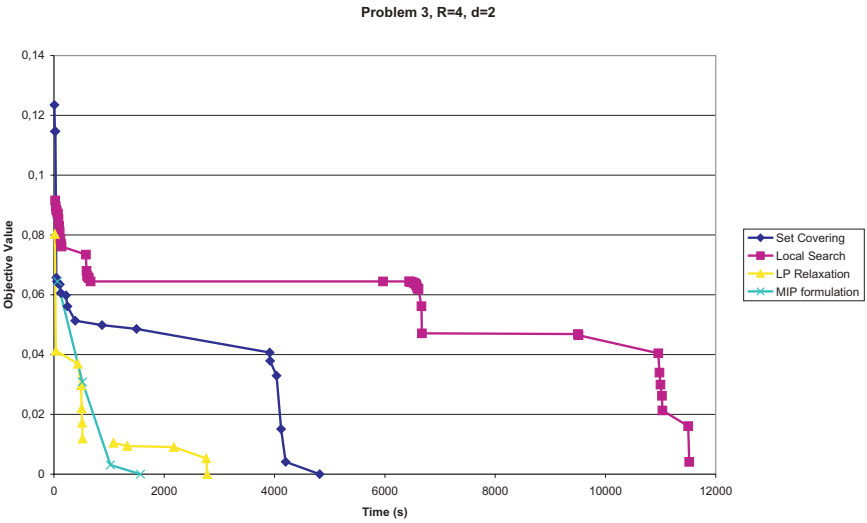

Fig. 5. Objective value versus computation time for problem 1

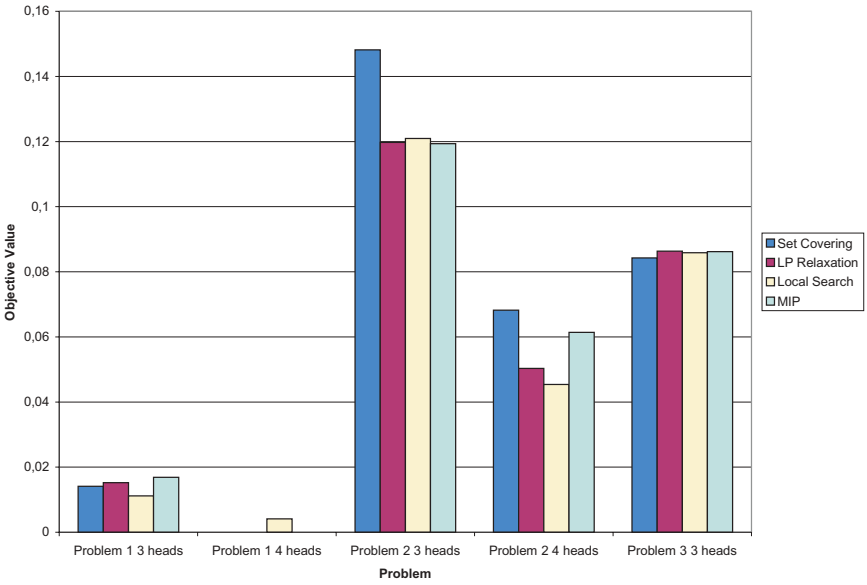

Fig. 6. Objective values for problems 1 to 3

5 Conclusion

In this paper we have presented a multiobjective mixed integer programming model for the problem of optimising intensity profiles and beam directions in intensity modulated radiotherapy planning. We have considered integrated and two phase

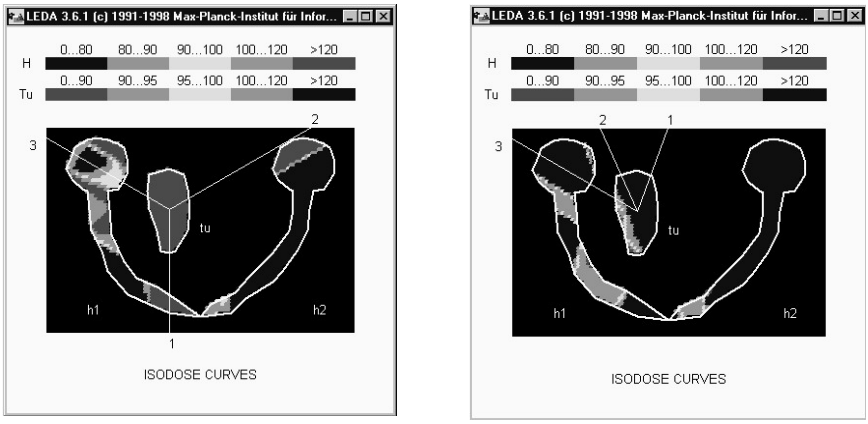

Fig. 7. Solution with optimised directions (right, objective value=0) versus solution with equidistant angles (left, objective value=0.4114)

methods for solving a weighted sum scalarisation of the model. The problem turns out to be hard when the number of beam directions R is small, and a feasible solution could often not be found within reasonable computation time. Therefore we considered three heuristic methods. The local search method changes each of the angles in steps of d degrees until a local optimum is encountered. Two phase methods try to determine good directions first, which are then fed into the intensity optimisation model. We proposed one method which uses the directions used in the optimal solution of the LP relaxation of the scalarised multiobjective model. A set covering method uses an optimisation model based on the intuitive idea of covering every voxel of the tumour and minimizing dose delivery to the organs at risk. Both methods are able to detect good directions. The numerical results indicate that optimisation of beam directions is very important in achieving good treatment plans. They also indicate that optimal plans are not always intuitive and thus illustrate that the intricate interdependencies between patient anatomy, beam geometry and intensity are difficult to capture for even experienced planners.

References

1. Küfer K-H, Scherrer A, Monz M, Alonso F, Trinkaus H, Bortfeld T, Thieke C (2003) Intensity-modulated radiotherapy – a large scale multi-criteria programming problem. OR Spectrum 25: 223–249
2. Das S, Cullip T, Tracton G, Chang S, Marks L, Anscher M, Rosenman J (2003) Beam orientation selection for intensity modulated radiation therapy based on target equivalent uniform dose maximization. International Journal of Radiation Oncology, Biology, Physics 55 (1): 215–224
3. Haas OCL, Burnham KJ, Mills JA (1998) Optimization of beam orientation in radiotherapy using planar geometry. Physics in Medicine and Biology 43: 2179–2193
4. Hamacher HW, Küfer KH (2002) Inverse radiation therapy planning – a multiple objective optimization approach. Discrete Applied Mathematics 118: 145–161

5. Holder A (2003) Designing radiotherapy plans with elastic constraints and interior point methods. Health Care Management Science 6 (1): 5–16
6. Holmes T, Mackie TR (1994) A comparison of three inverse treatment planning algorithms. Physics in Medicine and Biology 39: 91–106
7. Hristov DH, Fallone BG (1997) An active set algorithm for treatment planning optimization. Medical Physics 24: 91–106
8. Kolmo P, Tervo J, Lahtinen J (1998) The use of cimmino algorithm and continuous approximation for dose deposition kernel in the inverse problem of radiation treatment planning. Physics in Medicine and Biology 43: 2539–2554
9. Lee S, Cho PS, Marks RJ, Oh S (1997) Conformal radiotherapy computation by the method of alternating projections onto convex sets. Physics in Medicine and Biology 42: 1065–1086
10. Marzi S, Mattia M, Del Giudice P, Caccia B, Benassi M (2001) Optimization of intensity modulated radiation therapy: assessing the complexity of the problem. Annali dell'Istituto Superiori di Sanità 37 (2): 225–230
11. Pugachev A, Li GL, Boyer AL, Hancock SL, Le QT, Donaldson SS, Xing L (2001) Role of beam orientation optimization in intensity-modulated radiation therapy. International Journal of Radiation Oncology, Biology, Physics 50 (2): 551–560
12. Pugachev A, Xing L (2001) Pseudo beam's-eye-view as applied to beam orientation selection in intensity-modulated radiation therapy. International Journal of Radiation Oncology, Biology, Physics 51 (5): 1361–1370
13. Rowbottom CG, Nutting CM, Webb S (2001) Beam-orientation optimization of intensity-modulated radiotherapy: clinical application to parotid gland tumours. Radiotherapy and Oncology 59 (2): 169–177
14. Shepard DM, Ferris MC, Olivera GH, Mackie TR (1999) Optimizing the delivery of radiation therapy to cancer patients. SIAM Review 41 (4): 721–744
15. Spirou SV, Chui CS (1998) A gradient inverse planning algorithm with dose-volume constraints. Medical Physics 25: 321–333
16. Stein J, Mohan R, Wang XH, Bortfeld T, Wu Q, Preiser K, Ling CC, Schlegel W (1997) Number and orientations of beams in intensity-modulated radiation treatments. Medical Physics 24 (2): 149–160
17. Tervo J, Kolmonen P (2000) A model for the control of multileaf collimator and related inverse problem. Inverse Problems 16: 1875–1895
18. Thieke C, Bortfeld TR, Küfer KH (2002) Characterization of dose distributions through the max and mean dose concept. Acta Onclogica 41: 158–161
19. Ulmer W, Harder D (1995) A triple Gaussian pencil beam model for photon beam treatment planning. Zeitschrift für medizinische Physik 5: 25–30
20. Wang XH, Mohan R, Jackson A, Leibel SA, Fuks Z, Ling CC (1995) Optimization of intensity-modulated 3d conformal treatment plans based on biological indices. Radiotherapy and Oncology 37: 140–152
21. Webb S (1994) Optimizing the planning of intensity-modulated radiotherapy. Physics in Medicine and Biology 39: 2229–2246
22. Wu Q, Mohan R (2000) Algorithms and functionality of an intensity modulated radiotherapy optimization system. Medical Physics 27 (4): 701–711
23. Wunderling R (1996) Paralleler und Objektorientierter Simplex-Algorithmus. PhD thesis, Konrad-Zuse-Zentrum für Informationstechnik Berlin. ZIB technical report TR 96-09
24. Xiao Y, Censor Y, Michalski D, Galvin J (2003) The least-intensity feasible solution for aperture-based inverse planning in radiation therapy. Annals of Operations Research 119: 183–203