

Hybrid Optimisation Technique for Radiotherapy Treatment Planning

O. C. L. Haas*, K. J. Burnham*, J. A. Mills+

*Control Theory & Applications Centre, Coventry University, CV1 5FB Coventry UK.

Tel: 44 1203 838972 Fax: 44 1203 838052 E-mail: ctac@coventry.ac.uk

+Department of Clinical Physics and Bioengineering, Walsgrave Hospital, CV2 2DX Coventry U.K.

Tel: 44 1203 538937

Abstract

This paper presents an optimisation strategy, based on a hybrid multi-objective genetic algorithm, to automate the planning of traditional radiotherapy treatment. The optimisation algorithm is formulated around the concept of Pareto optimality to exploit a population based search procedure thereby considering each of the objectives independently. It explores the solutions belonging to the Pareto optimal set prior to concentrating on the best regions of the search space. Once the Pareto optimal set has been determined, *a posteriori* articulation of the objectives is used to provide the clinicians with a few alternative solutions, all of them non dominated and therefore mathematically optimal, leaving the final decision to the clinicians.

1. Introduction

The use of radiation to treat cancer patients is over one hundred years old [1]. However it is only recently, with the development of fast computers that the automation of radiotherapy treatment planning has become possible.

Traditionally, the planning of radiotherapy treatment is performed by human operators using methods based on trial and error, relying on experience and good practice [2]. The aim pursued by the treatment planners is to provide the clinicians with a plan that conforms well to the region to be treated whilst at the same time sparing the critical structures. This can be achieved by determining the optimal beam arrangement, i.e. number, orientation, shape and intensity modulation of the beams.

The aim of beam orientation is to combine beams such that they overlap over the diseased region and avoid critical structures. The use of beams incoming from several different directions leads to an increase in the relative dose delivered to the diseased region.

The shaping of the beams helps to focus radiation solely on the diseased region as seen from the beam. This can be achieved by making use of lead blocks, and/or multileaf collimators, to shield healthy tissues from radiation [2].

The determination of optimal beam intensity modulation is a much more complex problem which depends on the beam modulation device used over the course of the

treatment. Conventional radiotherapy makes use of wedge shaped compensators to modulate the beam intensity. The beam modulation resulting from the use of wedges depends on the wedge angles and the individual beam weightings [2].

Optimisation problems in radiotherapy have traditionally been solved efficiently using cost functions involving weighted sums, see for example [2, 3, 4]. However, such approaches rely on the skill of an operator to select the various weightings associated with the different objectives. As the selection of these objective weightings is performed manually there is no guarantee that they are optimal, nor is it true that a single solution found by the search algorithm is the best compromised solution achievable. Indeed, it has been observed in [4, 5] that a mathematical optimal solution found, whilst optimising beam weightings, could lead to unbalanced plans, which may not be clinically acceptable.

In contrast to the weighted sum approach, which reduces the multi objective problem into a single objective problem, the search technique presented in this paper enables the optimisation of various objectives in parallel. The optimisation is performed using a heuristic search technique, namely a multi-objective genetic algorithm. The approach is similar to that in [6] in that Pareto ranking forms the basis of the selection process. The approach has, however, been modified to enable the treatment planner to guide the optimisation as the search progresses towards a promising region, of the solution space, by making use of a decision maker. The best compromise solution can then be selected from the Pareto optimal set *a posteriori*, using goal attainment weighted sum of the objectives, or by considering the objectives by order of importance.

2. Preliminaries

The radiotherapy treatment planning problem may be visualised in two stages. The first stage, which is to establish the number and orientation of the beams, is achieved using the methods described in [5, 7]. This paper focuses attention on the second stage, that is of estimating/evaluating the optimum beam weight/wedge angle combination.

Wedges are the simplest beam modulation device, see Figure 1. They are differentiated according to their 'wedge angle' which is the slope of the isodose contour

at standard measurement depth (i.e. 10 cm depth) [2, 5]. In this paper it is assumed that use is made of motorised wedges which can produce any wedge angle between 0° and 60° [8].




Fig. 1: Illustrating 30° and 45° wedges that can be inserted into the head of the linear accelerator to modulate the beam intensity across the field into a wedge shape.

In order predict the dose distribution resulting from a set of wedges, the interactions between X-ray radiation and the various body structures are modelled, see [5]. In this paper, use is made of a matrix based divergent pencil beam model, which takes into account the effects of inhomogeneities, in-air-profile, penumbra and patient contour correction [2, 5]. The beam modelled, which is described in [5], relates the total dose distribution \mathbf{d} to a set of beam profiles \mathbf{b}_i via the following relationship

$$\mathbf{d} = \sum_{i=1}^{N_{BEAM}} \Phi_i \mathbf{b}_i \quad (1)$$

where the index $i=1 \dots N_{beams}$ denotes the beam considered, \mathbf{b}_i are the vectors describing the individual beam profiles and Φ_i is the dose calculation matrix details of which may be found in [7].

3. Objectives of radiotherapy treatment

The optimisation of radiotherapy treatment plans requires the formulation of the outcome of treatments into mathematical expressions. There are various methods currently available to predict the clinical acceptability of a radiotherapy treatment plan. Radiobiological models describe the effect of radiation at the cell level to predict chance of complication due to overdosage or probability of controlling tumours [3]. However the predicted outcome of treatment using radiobiological models is still a controversial issue [3]. In this work, a more traditional objective formulation based on prescribed and delivered dose is adopted.

Doses are expressed in terms of percentages of the dose prescribed at the isocentre, i.e. the centre of the cancerous region. The dose prescribed to the isocentre depends on the location, the extent, and the nature of the cancer. The dose limiting factors include the risk of narcosis [1, 2] and the presence of critical structures which tolerate only a small quantity of radiation.

In each radiotherapy treatment plan three regions of interest are considered: the planning target area (PTA), the organs at risk (OARs) and the other healthy tissues (OHT) which include all the body structures not included in the PTA and the OARs. Making use of a

vector notation, all grid points belonging to the same region of interest are grouped together and represented under a vector form such that \mathbf{d}_{PTA} , \mathbf{d}_{OAR} , \mathbf{d}_{OHT} , δ_{PTA} , δ_{OAR} and δ_{OHT} denote predicted and prescribed doses for the PTA, the OARs and the OHT respectively. The radiotherapy treatment planning objectives can then be expressed as follows:

$$C_{PTA} = [\delta_{PTA} - \mathbf{d}_{PTA}]^T \mathbf{W}_{PTA} [\delta_{PTA} - \mathbf{d}_{PTA}] \quad (2)$$

$$C_{OAR} = [\delta_{OAR} - \mathbf{d}_{OAR}]^T \mathbf{W}_{OAR} [\delta_{OAR} - \mathbf{d}_{OAR}] \quad (3)$$

$$C_{OHT} = [\delta_{OHT} - \mathbf{d}_{OHT}]^T \mathbf{W}_{OHT} [\delta_{OHT} - \mathbf{d}_{OHT}] \quad (4)$$

where C_{PTA} , C_{OAR} , C_{OHT} are individual costs relating to PTA, OARs and OHT; \mathbf{W}_{PTA} , \mathbf{W}_{OAR} , and \mathbf{W}_{OHT} are diagonal adaptive weighting matrices, which are fully described in [5, 9], for the PTA, the OARs and the OHT respectively. These matrices enable individual objective weightings to be assigned for the different regions of interest and also allows for constraints to be relaxed on dose points satisfying the dose requirements. A low C_{PTA} produces a high and uniform dose in the PTA, a low C_{OAR} produces a low dose in the OARs and C_{OHT} helps to keep the dose in the OHT low, leading to a uniform dose in the OHT.

4. The Multi-objective Approach

The multi-objective genetic algorithm (MOGA) described in this paper belongs to a family of evolutionary algorithms aimed at solving multi-objective optimisation problems. It is used to determine all the non dominated solutions, belonging to the Pareto optimal set [5, 10], prior to selecting the most suitable solution using *a posteriori* articulation of the objectives [5].

Genetic algorithms

Genetic algorithms (GAs) are a guided random search algorithms that are inspired from the natural principle of evolution. Essentially a population search based technique, they are particularly suited to multi objective optimisation problems. Another advantage of GAs is their robustness and problem independence [11, 12]. Further, as opposed to deterministic techniques such as gradient search methods, GAs have been shown to be capable of finding the global optimum in a search space including several local optima [6, 11].

Factors to take into consideration to maximise the performances of a GA include the coding system, the search operators, the selection process, the size of the population and the probabilities for the various operators to be used.

Coding strategy

Alternative coding strategies have been investigated. However, no significant advantage could be found in the use of integer or real coding over traditional binary coding. Therefore, a binary coded representation has been adopted to code the physical values of the beam weights and wedge angles. Seven bits are used to code the beam weights and the wedge angles, which can vary in the range [0 100%] and [-60° +60°] respectively.

Search operators

In this work use is made of both traditional genetic search operators, such as mutation and crossover, and more specialised operators that have been developed to solve the particular radiotherapy optimisation problems. Crossover combines pieces of the parents to produce an offspring and mutation selects randomly an element of the chromosome and changes its value. A specialised operator based on cautious least squares is used to generate a beam intensity and a wedge angle from a dose prescription for a particular beam [7, 13].

Selection mechanism

The selection of a solution for 'reproduction' is based on its quality, a measure of which is given by an objective function or cost function. The value of the objective function is not used directly but transformed into a fitness function which indicates a probability of selection. Such a probabilistic process is one of the central features of GAs. It introduces an element of 'chance' into the selection process. The selection mechanism presented in this paper is based on Pareto-ranking [5, 6]. In contrast to traditional ranking techniques, where all the solutions are ranked according to a single objective, with each rank being unique, Pareto ranking, assigns to all the non dominated solutions an identical rank of 1, see Figure 2. The dominated solutions being differentiated by the number of solutions by which they are dominated, see Figure 2. The rank of each possible solution is calculated as follows

$$r = 1 + N_{\text{dominated}} \quad (5)$$

where r is the rank and $N_{\text{dominated}}$ is the number of solutions by which the solution considered is dominated.




Figure 2 Illustrating the concept of Pareto optimality and non dominated solutions for two objectives.

This approach makes use of the concept of dominance which can be explained as follows. Considering a minimisation problem, in which the i^{th} solution is represented by an objective function vector (or criterion vector) $\mathbf{v}_i = [v_{i1} \ v_{i2} \ \dots \ v_{in}]$, such a solution is said to be non dominated if there is no element in \mathbf{v}_i which is greater than the corresponding element in any other objective function vector \mathbf{v}_j and that at least one element in \mathbf{v}_i is less than its corresponding element in \mathbf{v}_j [10].

Pareto ranking could lead, after a few generations, to a population composed solely of non dominated solutions, with an identical probability of being selected. At this

stage, the selection of a new individual would become a random process. However, in most problems, solutions located at one extreme of the Pareto set may be very different to solutions located at the other. It is generally accepted that a combination of two such solutions is not likely to give rise to a better solution [6, 11, 12]. To differentiate individual solutions of similar rank, and favour the combination of individuals that are more likely to produce improved solutions use is made of the so called progressive articulation of the objectives. This can be implemented using a weighted sum of the various objectives. For each solution defining $\tilde{r}_j = r_j + \Delta r_j$, where Δr_j is given by

$$\Delta r_j = \alpha \sum_{i=1}^n \left(\lambda_i \frac{v_{ij}}{\max(\mathbf{v}_j)} \right) / \sum_{i=1}^n \lambda_i \quad (6)$$

in which λ_i , v_{ij} and $\max(\mathbf{v}_j)$, $i = 1 \dots n$, $j = 1 \dots m$, are, respectively, the objective weightings, the values of the j^{th} objective in the i^{th} solution and the maximum value of the j^{th} objective in all m solutions. The scalar quantity $0 \leq \alpha < 1$ is chosen to emphasise the relative importance of the rank r and the weighted sum of the objectives.

Δr can be used to investigate the entire Pareto optimal set (when no *a priori* information is available) or to focus the search on a particular region of the Pareto optimal set. The values λ_i may vary as the search progresses in a deterministic or in a heuristic manner by:

- Alternating the objectives randomly i.e. setting one of the weightings associated with one objective to unity whilst the other weightings are set to zero;
- Alternating the objectives sequentially in an orderly manner;
- Combining the objective weightings randomly, such that the weighted sum of the objectives evolves in a random manner;
- Varying the objective weightings around a fixed set of values.

Such an approach is used at the beginning of the search to explore different regions of the solution space prior to concentrating on the 'best', in some sense, region.

Selecting acceptable solutions from the Pareto optimal set

Having determined the Pareto optimal set it is necessary to choose the best compromised solution from all the non-dominated solutions. Ultimately, the clinicians are responsible for the final decision, therefore the aim of the decision maker, which is used to select acceptable solutions, is to present the clinician with a limited number of solutions to choose from. Provided that all the criteria set by the clinicians can be satisfied a single solution may be selected by the decision maker. In most cases, however, all the objectives cannot be achieved simultaneously, therefore compromised solutions which favour different objectives may have to be presented to the clinician. The clinician may then make a decision according to some specific clinical knowledge not taken into account in the optimisation process. The most common approaches for selecting suitable solutions make use of:

- coefficients to weight the relative merits of the objectives,

- measures to indicate the degree to which each of the objectives is attained,
- set priorities to consider the objectives according to their relative importance.

In order to illustrate the need to present a clinician with more than one solution in cases where a compromise solution has to be selected, consider a plan with two OARs close to the PTA, such that it is not possible to deliver a low dose to both OARs when the dose in the PTA is within prescription. In this case, different combinations of objective weightings may be selected to favour the PTA and one of the OARs or to favour both OARs and reduce the dose delivered to the PTA region. These two solutions may then be presented to a clinician who will use clinical knowledge which may not have been considered in the optimisation process to select a solution which is most appropriate to a particular patient's needs.

A disadvantage of objective weighting is that the weights depend on the mathematical expression used to calculate the objectives for the different regions, which makes them unique to a particular optimisation software. Clinicians are not usually consulted to assess a plan in terms of weightings associated with various objectives, but they may be more adept to handle objectives expressed in terms of goals or percentages. For example, considering the PTA region, it is possible to prescribe that 100% of the cells within the PTA receive a dose d_{PTA} within $\pm 5\%$ of that prescribed. However, it may be advantageous to allow 1% of the PTA region to receive a dose less than specified if the maximum dose received by an OAR can be reduced by 10%.

The use of priorities can also be beneficial where the clinician could specify that the dose in the PTA should be within 5% of that prescribed. Then, provided that this can be achieved a subsequent aim would be to reduce the dose to some OAR, to a specific level, and finally, if possible, minimise the extent of hot spots in the OHT.

Algorithmic representation

A general algorithmic representation of the MOGA used in this work is given as follows:

- 1) Create an initial population by randomly generating solutions and seeding the population with some good solutions.
- 2) Calculate the value for each objective C_{PTA} , C_{OAR} , C_{OHT} .
- 3) Deduce the modified rank \tilde{r} .
- 4) Deduce the fitness values from \tilde{r} .
- 5) Select parents from the population according to their fitness to produce offspring.
- 6) Combine each pair of parents together to produce a pair of offspring or create a new offspring from a single parent.
- 7) Evaluate the objective costs for the new solutions.
- 8) Deduce the modified rank \tilde{r} .
- 9) Deduce the fitness values from the value of \tilde{r} .
- 10) Delete members of the population to insert all the non dominated solutions (delete all dominated solutions provided that the number of individuals in the population is not smaller than the original population size, otherwise delete only the worse individuals).
- 11) If the maximum number of iterations is reached or the population has converged then: use the *a posteriori* decision maker to select acceptable

solutions from the Pareto optimal set, (otherwise go back to step 5).

5. Optimisation of wedges

To illustrate the approach, consideration is given to a typical computed tomography slice comprising three regions of interest, namely the PTA, the OARs (including the spine and the rectum) and the OHT (including the two femoral heads as well as the other structures not included in the PTA and OARs).

The MOGA is used to determine the Pareto optimal set for three objectives, namely C_{PTA} , C_{OAR} and C_{OHT} associated with the PTA, the OARs and the OHT respectively. It is observed that for a given beam orientation, most solutions offer similar acceptable cost for C_{PTA} . As a consequence, the selection of the best compromise solution uses the other two objectives C_{OAR} and C_{OHT} see Figure 3. It is found that a decrease in the number of objectives leads to a decrease in the number of non dominated solutions, i.e. previously non dominated solutions are now dominated.




Figure 3: Illustrating the selection of two candidate solutions lying on the Pareto optimal set C_{OAR} and C_{OHT} .

Two solutions were selected from the Pareto optimal set of Figure 3. The first solution, represented by the point marked 'x', for which the isodose plot is shown in Figure 4, is chosen to produce a relatively low error in the OARs, whilst keeping the error in the OHT at a reasonable level. However, the isodose plot reveals that although the sparing of the OAR is acceptable, the plan is unbalanced, with the right side of patient (left side in Figure 4) receiving a dose some 10% higher than the left side.




Figure 4 Solution chosen for its relatively low error in the rectum resulting in an unbalanced plan.

In order to compensate for this effect of creating a hot spot, a new solution was chosen, indicated by the point marked 'o' in Figure 3, with the corresponding isodose plot, shown in Figure 5. This solution has an equivalent normalised cost for both the OAR and OHT.




Figure 5 Isodose contour plot for a solution chosen to deliver a low dose to the OHT resulting in the equalisation of the dose given to the femoral heads.

It can be observed that the hot spot on the left side of the plan, c.f. Figures 4 and 5, has been removed, at the expense, however, of a higher dose delivered to the OARs; with two thirds of the OARs receiving a dose of at least 30%. Such a dose level is, however, still acceptable. Further, the dose homogeneity in the PTA has been improved.

It is interesting to note that the use of a Pareto optimal approach overcomes the difficulties reported in [4] where the optimised plans were relatively likely to be unbalanced with unacceptable hot spots. In particular there is no need to re-start the optimisation algorithm with a new set of objective weightings. Since all the solutions are optimal, a more appropriate solution can be determined by changing the preference of the decision maker *a posteriori*. Alternatively, a graphical method can be used to select a solution from the Pareto optimal set, although the objectives have not been specifically aimed at reducing or equalising hot spots. Another advantage of the Pareto optimal formulation is that it enables the degree by which the mathematical objectives are correlated to be determined, thus leading to a reduced number of objectives considered [6].

6. Conclusion

This paper has shown that a multi-objective genetic algorithm with Pareto ranking can be used to advantage when optimising traditional radiotherapy treatment plans. The approach, which has been evaluated on a realistic test case, is able to overcome problems encountered by traditional approaches which combine all the objectives into a single weighted sum.

A problem with the proposed approach, however, is the need to evaluate a large number of candidate solutions, which makes it a comparatively slow process and impractical for current use within a radiotherapy clinic. Whilst this is likely to be less of a problem in the future, with computer power becoming more widely available, alternative techniques are presently being pursued in an attempt to overcome this immediate problem.

By exploiting elements of similarity between types of cancer, e.g. anatomical location and patient geometry, together with an abundance of data, it is possible to make use of an artificial neural network (ANN) to speed up the search for an optimal solution. The ANN would be trained using typical treatment plans which are appropriate for a range of patients. The ability of an ANN to quickly determine initial solutions relevant to a particular patient would make treatment requiring complex planning procedures a practical clinical reality.

References

- [1] Mould R.F. A Century of X-rays and Radioactivity in Medicine. Institute of Physics Publishing, Bristol and Philadelphia, 1993
- [2] Webb S., *The physics of three-dimensional radiation therapy: conformal therapy, radiosurgery and treatment planning*, IOPP, Bristol, 1993
- [3] Smith A.E., *Radiation therapy physics*, Springer Verlag : Berlin, 1995
- [4] Oldham M. Webb S. A comparison of conventional 'forward planning' with inverse planning for 3D conformal radiotherapy of the prostate, *Rad. Oncol.*, 248-262, 1995
- [5] Haas O.C.L., *Optimisation and Control Systems Modelling in Radiotherapy Treatment Planning* PhD Thesis, Coventry University UK, 1997
- [6] Fonseca C. *Multiobjective Genetic Algorithms with Application to Control Engineering Problems*, PhD Thesis, The University of Sheffield, UK, 1995
- [7] Haas O.C.L., Burnham K.J., Mills J.A., Optimisation of beam orientation in radiotherapy using planar geometry, in press *Physics in Medicine and Biology*, 1998
- [8] Beavis A.W., Weston S.J., Whitton V.J., Enhanced dynamic wedge - the coming of age of intensity modulated conformal therapy planning. *Radiology UK* 96, *The British Journal of Radiology*, Suppl to vol. 69, 99, 1996
- [9] Haas O.C.L., Burnham K.J., Mills J.A., On Improving physical selectivity in the treatment of cancer: a systems modelling and optimisation approach, *Control Engineering Practice* 5(12), pp. 1739-1745, 1997
- [10] Steuer R.E., *Multiple criteria optimization: Theory, Computation, and Application*, Wiley & Sons, USA, 1986
- [11] Beasley D., Bull D.R., Martin R.R., 'An overview of genetic algorithms: part1, fundamentals, *University Computing*, 15 (2) 58-69, 1993
- [12] Reeves C R. *Modern Heuristic Techniques for Combinatorial Problems*. Blackwell Scientific Publications, 1993
- [13] Haas O.C.L. Burnham K.J. Mills J.A., Adaptive error weighting scheme to solve the inverse problem in radiotherapy *Proc. 12th Int. Conf. on Systems Engineering, ICSE'97 Coventry University UK* pp. 290-295, 1997
- [14] O.C.L. Haas, K.J. Burnham, M.H. Fisher, J.A. Mills, "Implementation of hot spot correction schemes for use in optimisation algorithms for conformation therapy", *Radiology UK* 96, *British Journal of Radiology*, Supplement vol. 69,100, 1996.