

Evolutionary approach to inverse planning in coplanar radiotherapy

V. Bevilacqua*, G. Mastronardi, G. Piscopo

Department of Electrical and Electronic Engineering, Polytechnic of Bari, Via Orabona, 4, 70125 Bari, Italy

Received 26 April 2004; received in revised form 27 June 2005; accepted 31 January 2006

Abstract

A unified evolutionary approach to coplanar radiotherapy inverse planning is proposed. It consists of a genetic algorithm-based framework that solves with little modification treatment planning for three different kinds of radiation therapy: conformal, aperture-based and intensity modulated. Thanks to evolutionary optimisation techniques we have been able to search for full beam configurations, that is beam intensity, beam shape and especially beam orientation. Unlike some previous works found in literature, our proposed solution automatically determines exact beam angles not relying solely on a geometrical basis but involving beam intensity profiles, thus considering the effective delivered dose. Our dose distribution model has been validated through comparison with commercial system: fixed the same beam configuration, both calculated beam shapes and DVH have been compared. Then we have tested the optimisation algorithm with real clinical cases: these involved both simple (convex target, far organs at risk) and complex (concave target, close organs at risk) ones. As stated by physician and by simulation with the same commercial system, our tools found good solutions in both cases using corresponding correct therapy.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Genetic algorithms; Conformal; Aperture-based; Intensity modulated radiation therapy

1. Introduction

Introduction radiation therapy treatment of cancers is typically achieved using X-ray beams directed at the tumour from different angles. In particular, external beam radiation therapy is used to treat various kinds of tumours. The radiotherapy equipment consists of a table—where the patient lies—and a linear accelerator mounted on a gantry, which can rotate around the table. Gantry angle and beam intensity can be set, as for the beam shape thanks to the multi-leaf collimator. The treatment goal is to deliver a high dose of radiation to the diseased tissue (target) while partially or totally avoiding normal tissue and organs at risk (OARs). There are different kinds of radiation therapy techniques. In our work, we have considered three among the most used: conformal radiation therapy (3D-CRT), aperture-based radiation therapy (ABRT) and intensity modulated radiation therapy (IMRT). Their main difference lies in beam intensity profile: constant for 3D-CRT, piecewise constant over discrete levels for ABRT, almost continuous for IMRT.

Forward planning is the manual process of finding the best beams configuration with respect to the specialist prescribed

dose. The operator decides some configuration parameters based on his experience, then executes a simulation to see the resulting dose distribution and repeats the process until some satisfactory treatment plan is found.

Inverse planning is the automated computation of the radiation therapy treatment plan that satisfies the prescribed dose over the target and over all critical structures (OARs with/out normal tissue). Usually, to achieve this result an optimisation algorithm is used to find the best solution, which is classified in terms of an objective function (or cost function). This describes how good (or bad) is a proposed solution with respect to the inverse planning problem.

Many methods can be found in literature to solve different kinds of optimisation problems for different kinds of radiation therapies [1]. Evolutionary techniques are a possible approach to the optimum solution search, but only a few examples exist about their application to inverse treatment planning [2–5,14]. In [4], evolutionary algorithms have been used as a better alternative to error back-propagation for evolving an artificial neural network (ANN). This ANN has been trained and then used to solve dose optimisation problem in inverse planning, with beams number and positions fixed for known typical cases. Nevertheless, thanks to genetic algorithms—a particular implementation of evolutionary algorithms—we have been able to find the optimal beam configuration, including beam orientation angles.

It must be noticed that, to our best knowledge, only a few recent articles give a proposal solution for the angles

* Corresponding author. Tel.: +39 080 596 3252; fax: +39 080 596 3410.
E-mail address: bevilacqua@poliba.it (V. Bevilacqua).

optimisation [5–8]. In [5–7,15], such a solution is found relying solely on a geometrical basis, not involving simultaneous beam weights and shapes determination, which is done in a second phase. In [8], hints on preferred beam directions to choose are given through a ‘beam selection’ metric.

According to the physician judgment and to the executed simulations on a forward planning commercial system,¹ the results obtained for some real clinical cases are suitable for effective radiation treatments. These results show how 3D-CRT can be used for simple clinical situations but how it fails for more complex cases, when ABRT and IMRT are needed. All three inverse planning problems have been solved using the same evolutionary-based framework while, in literature, different algorithms are proposed for different kinds of therapies.

2. Algorithm proposed

The designed software algorithm is based on two main procedures: the dose distribution model and the optimisation algorithm itself. The former takes into account the way radiation propagates through the volume. This happens independently of intensity initial value and/or intensity profile at the skin surface, so the model is actually the same for every kind of radiation therapy considered. The latter is the actual implementation of genetic algorithm and varies from therapy to therapy, although they share the same structure and the same cost-function evaluation.

2.1. Dose distribution model

Implemented model is based on single ray tracing from one fixed direction. Before the projection/retro-projection a rotation from generic angle to the fixed direction must occur. After retro-projection only, an inverse rotation is needed to obtain correct dose distribution over each section. As directions are coplanar, an axis of symmetry exists (z) and all rotations are intended around this axis (in xy -plane). Absorption of radiation through body volume is taken into account multiplying intensity by a function of depth along the path. This function approximates an experimental absorption curve given by the physician. It takes care of lateral beam diffraction in a simple way, having a maximum just after skin surface and then exponentially decaying. In this first approach to inverse planning problem horizontal beam divergence is not modelled, nor is vertical beam divergence. In such context, the isocenter position is a trivial matter, so it has been fixed at the centre of given volume. Easy modifications are needed to change this feature and let isocenter position free to be set.

2.2. Optimisation

We have chosen genetic algorithms as optimisation technique mainly because of the complexity of inverse

planning problem. First of all, their stochastic nature permits to investigate fast enough for enough possible solutions in a huge search space, as can be the one determined by the variables: beam angles, beam intensity and beam shapes. Second, these techniques are suitable for problems with a global optimum and multiple local optimal solutions.

This is due especially when beam angles are also optimised, as it has been stressed that in such situation many local maxima (minima) appear in objective (cost) function [9,10]. Genetic algorithms [12,13] emulate natural evolution of species, evolving a population according to the principles of natural selection and survival of the fittest. Every individual in a population represents a possible solution to the given problem. Every one of them receives a score (penalty) according to how good (bad) it is as a solution to the problem.

The best (or fittest) ones are given more chances to reproduce themselves and become parents for the new generation. Every couple of parents gives two descendants (offsprings) that share some features from each one of them, possibly becoming even fitter than them. In this way a new generation is created from the previous one, having more and more good features spread in the population. If genetic algorithm has been well designed, the population will converge to the best individual—that is, the global optimal solution to the problem.

For our purposes, every individual must represent a possible treatment plan. The genes that define an individual are the parameters that define a plan, that is, beam configuration parameters. This part of implementation changes accordingly to the particular radiation therapy considered, as change the parameters that define some treatment plan. Anyway, they must comprise:

- beam direction
- beam aperture for each vertical section
- beam intensity profile for each vertical section.

Especially, the last one changes from therapy to therapy: a single scalar value in 3D-CRT or a function for each section in ABRT and IMRT. All of these parameters must be known for every beam that belongs to the treatment plan. The number of beams used in the treatment is fixed for all individuals.

Starting from genes, the dose distribution model enables to calculate the dose distribution associated with some plan. Then a specially designed cost function gives the penalty for the individual: this represents the distance of the realized distribution dose from the ideal one prescribed by the physician.

Prescriptions limits change for every noticeable structure and also change the rules in evaluating the distance for each one of these. Guided by the physician knowledge, we have classified six different kinds of tissue with their respective dose limits (dx) and penalty rules (p) given the delivered dose (d) over a region voxel. Rules are shown in formulas below and, for some tissues, in Figs. 1 and 2:

- ‘Target’ (T)

¹ PLATO complete System, <http://www.nucletron.com>

$$p = \begin{cases} c_{T1}(d_{T1} - d) + s_{T1} & \text{if } d < d_{T1} \\ 0 & \text{if } d_{T1} \leq d \leq d_{T2} \\ c_{T2}(d - d_{T2}) + s_{T2} & \text{if } d_{T2} < d \leq d_{Tmax} \\ \infty & \text{if } d > d_{Tmax} \end{cases}$$

- ‘Partially critical’ (PC)

$$p = \begin{cases} c_p d & \text{if } d \leq d_p \\ \infty & \text{if } d > d_p \end{cases}$$

- ‘Absolutely critical’ (AC)

$$p = \begin{cases} 0 & \text{if } d = 0 \\ \infty & \text{if } d > 0 \end{cases}$$

- ‘Minimally critical’ (MC)

$$p = c_M d$$

- ‘Body’ (B)

$$p = c_B d$$

- ‘Outside’ (O)

$p=0$ (do not care).

The cost of a single plan is obtained by summing the average penalty over each kind of region: it is a typical implementation of dose-based cost. This part must be carefully designed, as cost function guides genetic algorithm in evolving the population and little mistakes lead to totally wrong results. Averaging only over exceeding voxels instead of averaging over all voxels with respect to some tissue is an example of such a mistake. Cost functions must be minimised by optimisation algorithms, so higher (lower) fitness is assigned to individuals with less (more) cost in the population.

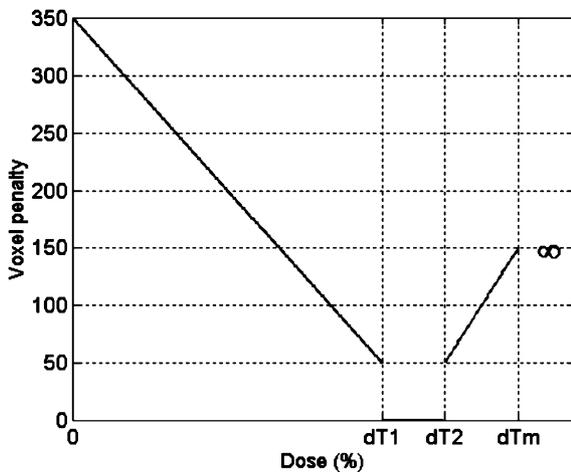


Fig. 1. Target penalty rule.

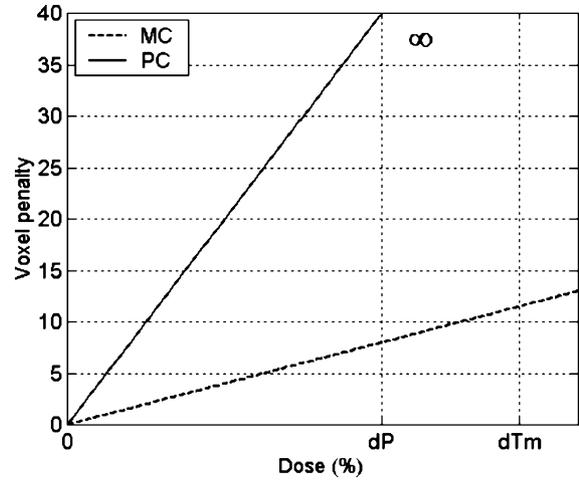


Fig. 2. OARs (MC, PC) penalty rules.

The genotype is directly represented as a real number and a roulette wheel selection algorithm has been adopted. Moreover, an uniform crossover comes: genes from two parents are mutually exchanged with some probability so that two new individuals are created.

To keep congruent treatment plans, exchanging a gene between two plans implies totally exchanging two beams between them—that is, exchanging all parameters related to a single beam. After crossover comes mutation, i.e. a single gene in every descendant can be randomly changed with low probability.

Depending on radiation therapy, the gene could be a beam angle or a beam weight. Newly created offsprings make up the new population for the next generation.

Evolution is obtained repeating this process, until some convergence criterion is reached. In this case, we have fixed a maximum number of generations, which seems to be valid for all three kinds of therapies and for all tests considered. Designed genetic algorithm implements elitism [11], i.e. at least one copy of the best individual from a generation passes unchanged to the next. The algorithm is also adaptive, i.e. mutation probability is increased when population becomes more homogeneous, and vice versa. This tries to find alternative solutions in search space even when population is converging to some optimum, giving the chance to still explore good areas of the solutions space.

Fig. 3 shows an example of a population evolution by means of cost function values: as it can be seen, having represented a particular test case extracted from a number of carried out runs, with a typical behaviour obtained on average, after one half of the run the algorithm finds the best individual and then makes population converge towards it.

Different simulations have been made on single 2D slices and full 3D volumes in order to set-up genetic algorithm parameters. Starting from classical values found within many articles and tutorials in literature, we have changed them to ensure convergence towards optimum and to obtain better performances. In the end, we have found that this evolutionary approach is robust enough to be quite insensitive with respect

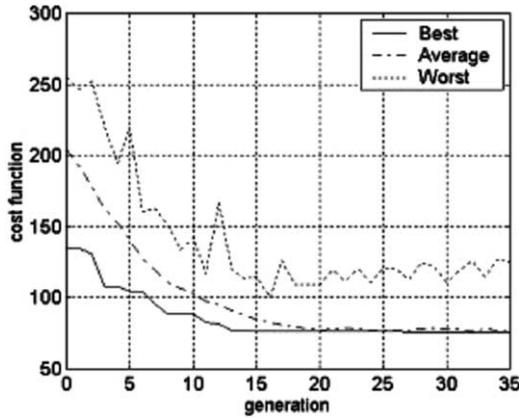


Fig. 3. Evolution of population.

to parameters set-up. Only mutation probability has been set adaptively to quite high values (1–10%), but this is used to act on a single gene of individuals instead of all genes.

3. Results

We have tested our algorithm with real clinical cases. These involved both simple (i.e. convex target, far OARs) and complex (i.e. concave target, close OARs) ones, as it can be seen in Figs. 4 and 5, respectively. It must be noticed that tests comprise full volumes and not only these single 2D images. Fixed dose limits are: $dP=80\%$, $dT1=94\%$, $dT=100\%$, $dT2=106\%$, $dTmax=115\%$. Dose is not expressed directly in Gy units but in terms of percentage of the dose prescribed at the target (100%). Fixed beam number is five for all kinds of therapy. Fixed genetic algorithm parameters are: population size 100, number of generations 35, crossover probability 80%, initial mutation probability 1%.

Among other graphs, the software produces a diagram widely used in clinical practice: dose–volume histogram (DVH). Given a tissue and its corresponding curve in this graph, if point (d,V) lies on the curve then $V\%$ of tissue volume receives a dose equal to or greater than $d\%$. DVH of a perfect dose distribution would have for the target a step function with corner $(100\%,100\%)$ and for other tissues some curves close to y-axis. Because of its use in the

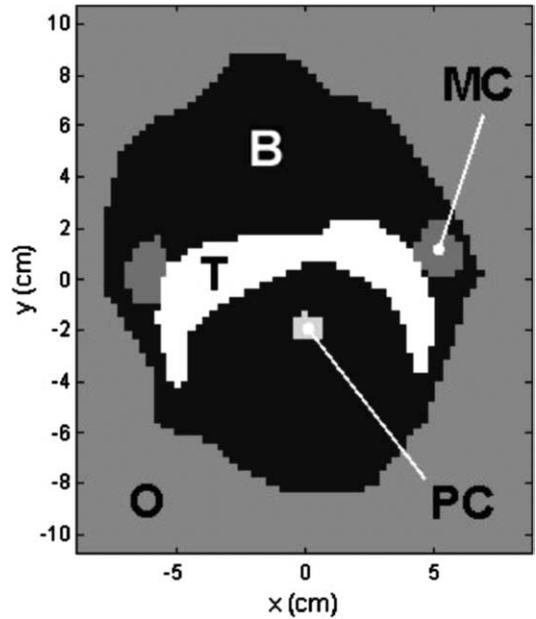


Fig. 5. Sample slice of complex clinical case.

real world, this diagram seems to us the best way to evaluate the quality of proposed treatment plans.

Figs. 6 and 7 show DVH for simple and complex cases respectively, both treated and resolved with 3D-CRT. These confirm that such a therapy is suitable for simple situations, and accordingly the software is able to find a good solution. As it can be seen, about 50% of target volume receives a dose in acceptable interval $[94\%,106\%]$. 3D-CRT fails instead for complex situations, and the software is not able to find a good solution too. This can be noticed from the ‘Target’ and ‘Par-Crit’ curves almost overlapped. They receive the same radiation, so dose on the former is limited by constraints on the latter. In Fig. 6, the curve relative to ‘Abs-Crit’ is hard to see as it lies on y-axis, thus confirming the absence of radiation over absolutely critical OAR.

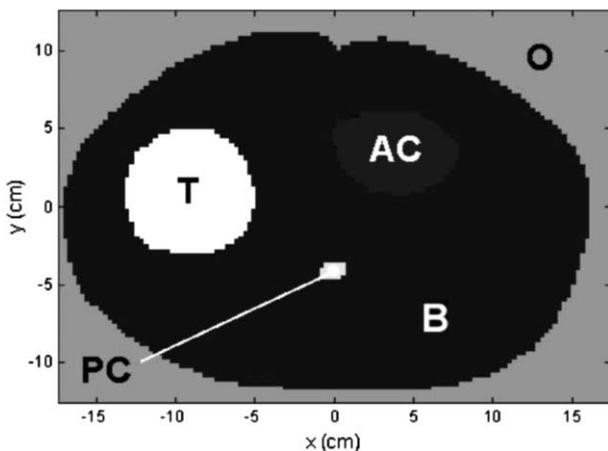


Fig. 4. Sample slice of simple clinical case.

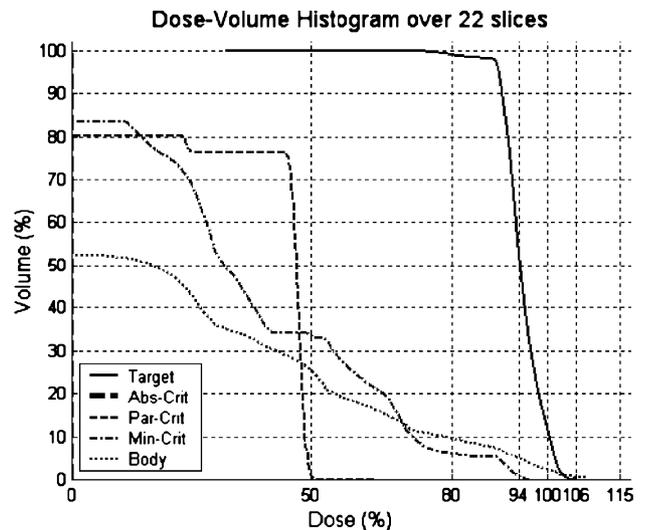


Fig. 6. DVH for 3D-CRT, simple case.

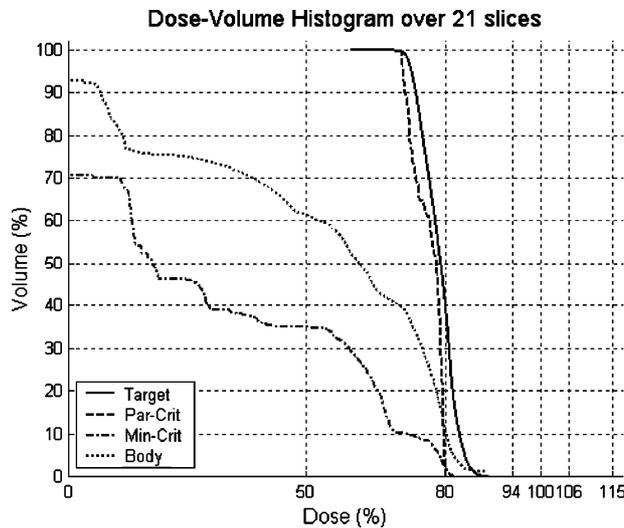


Fig. 7. DVH for 3D-CRT, complex case.

Some interesting features from DVHs in Figs. 6 and 7 are shown in Tables 1 and 2, respectively. In the first data column they report ‘correct volume’: this is the volume fraction of tissue that receives correct dose—that is a dose respecting ideal constraints—when this makes sense.

Figs. 8 and 9 show DVH for complex case only, treated and resolved with ABRT and IMRT, respectively. It must be noticed that these graphs refer to runs on a single 2D slice (Fig. 5) in order to show how the software is able to work with this target geometry. Anyway, the algorithm has been tested with full volumes, when entire tissue shapes are taken into account and averaging over several slices makes this capability less noticeable. These DVHs confirm that such therapies are suitable for complex cases: accordingly, the software is able to find a dose distribution that ‘follows’ complex target shape. With ABRT almost 75% of target volume receives acceptable dose and then this steeply decreases. With IMRT the fraction raises to 85% and dose decreases more slowly. Tables 3 and 4 report quantitative data regarding the same therapies applied to the same complex case, but this time relative to full volumes. Volume fraction of target hit correctly is the same as for the first example, 3D-CRT with simple case.

Delivered dose over subsequent slices is shown in Fig. 10 by means of several 2D splash-diagrams. The splash-diagrams are typical diagrams used by the physicians to represent the goodness of the obtained results in terms of grey tones. In fact with five beams, whose angle is shown by using the \times symbol in the previous Fig. 10, it is possible to argue that the grey tones

Table 1
Dose coverage for 3D-CRT, simple case

Tissue	Correct volume (%)	Avg. dose (%)	Max. dose (%)	Min. dose (%)
T	53	94.7	106.0	32.2
AC	100	0.0	0.0	0.0
PC	100	37.3	63.8	0.0
MC	/	37.4	96.3	0.0
B	/	26.3	108.0	0.0

Table 2
Dose coverage for 3D-CRT, complex case

Tissue	Correct volume (%)	Avg. dose (%)	Max. dose (%)	Min. dose (%)
T	0	78.5	89.0	59.9
PC	100	76.3	80.0	70.5
MC	/	30.8	81.6	0.0
B	/	50.9	88.3	0.0

represent the amount of energy delivered in every zone, and that the grey is lighter where more than one ray meet and overlap. They are useful to prove how the calculated beam aperture ‘follows’ target extension for each section and for each point of view, resulting in beam eye-view when these apertures are seen stacked in vertical direction (BEV, Fig. 11). Beam shapes determined by the software have been successfully compared with those computed by the forward planning commercial system. Fixed the same beam angle set they give the same beam shapes, thus confirming the correctness of the system geometry and ray projection technique. With the same beam configuration, dose distribution has been generated for both commercial system and our software: Figs. 12 and 13 show the relative DVHs (again, in our diagram AC curve lies on y-axis).

In Fig. 14, an example dose distribution over a single slice is presented in a 3D fashion, with height corresponding to delivered dose. This can be compared with prescribed dose over the same slice (Fig. 15), where height represents ideal values to obtain (over the target) or limits to respect (over critical structures). A typical solution is represented by diagram in Fig. 16, where arrow directions give beam orientations and arrow lengths give beam weights (actually beams ‘fire’ from outer border towards the centre, and not the opposite like arrows do). As it can be seen, the algorithm can indicate effective number of beams required by assigning negligible intensities to unneeded ones.

Cost function can be altered easily to give different importance to different features required for dose distribution,

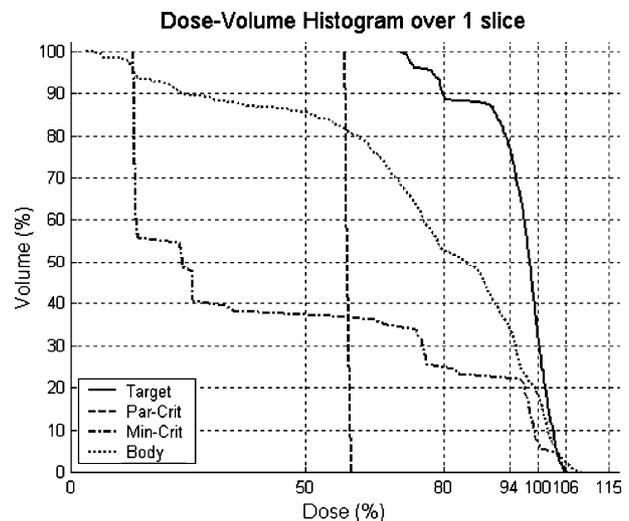


Fig. 8. DVH for ABRT, complex case.

Table 3
Dose coverage for ABRT, simple case

Tissue	Correct volume (%)	Avg. dose (%)	Max. dose (%)	Min. dose (%)
T	57	92.3	106.0	48.9
PC	100	41.1	72.0	7.9
MC	/	69.9	104.4	25.5
B	/	30.7	105.7	0.0

Table 4
Dose coverage for IMRT, simple case

Tissue	Correct volume (%)	Avg. dose (%)	Max. dose (%)	Min. dose (%)
T	53	91.3	105.0	54.7
PC	100	66.7	80.0	54.8
MC	/	32.2	95.7	0.0
B	/	55.1	110.3	0.0

such as homogeneity, hot and cold spots. Results are given in the same way and with the same tools which physicians are used to: dose distribution over 2D sections, BEV, DVH with resuming tables.

Actual implementation has been developed with Matlab routines, but only basic element-by-element arithmetic operations and sub-matrix indexing occur.

Thus, all functions can be easily ported to some specific programming language, such as C/C++, and to some specific operating system. On a 600 MHz-AMD PC it took about 30 min to run optimisation algorithm for a complete volume, so promising performances could be obtained with faster implementation running on faster machine like the ones commonly used in radiation therapy treatment planning.

Approximations have been introduced only in dose distribution model, regarding various kinds of divergences and other effects of radiation propagation. Simple modifications are required in order to improve software and make it closer to standard simulation systems (divergence, electron transport, vacuum absorption). Our dose distribution model has

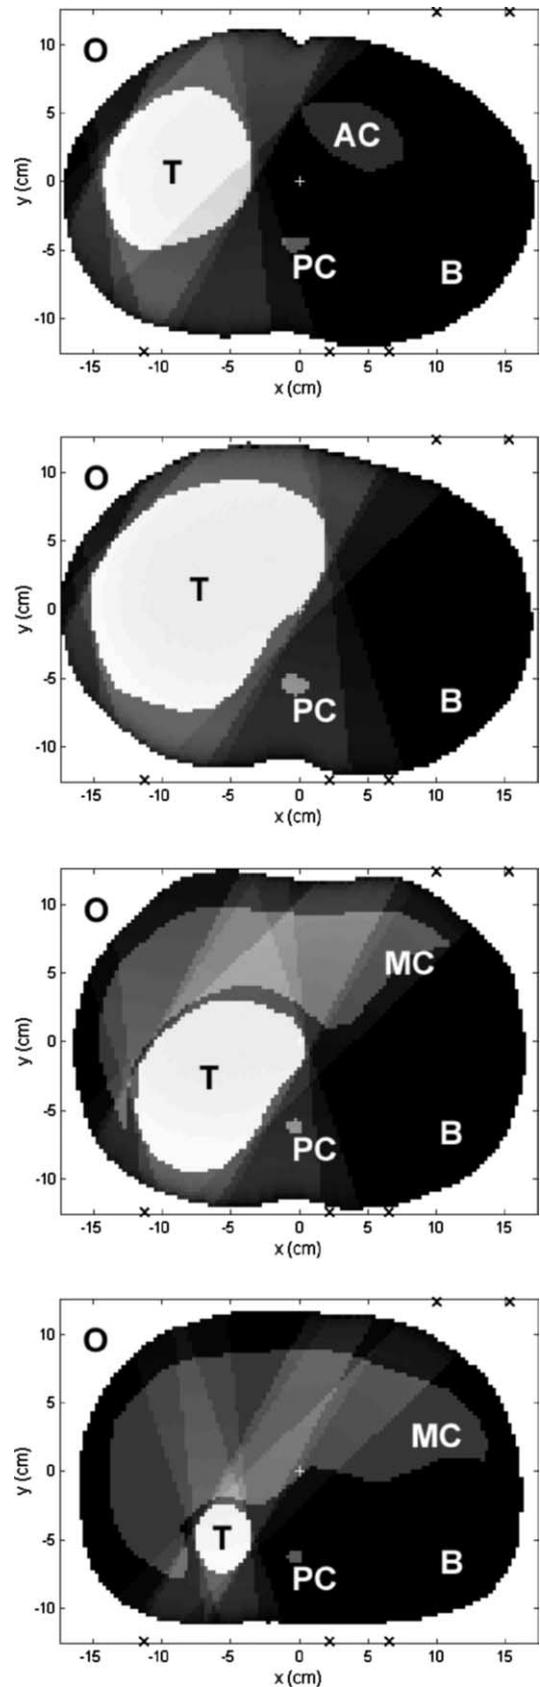


Fig. 10. Delivered dose overlapped to prescribed dose for subsequent 2D slices.

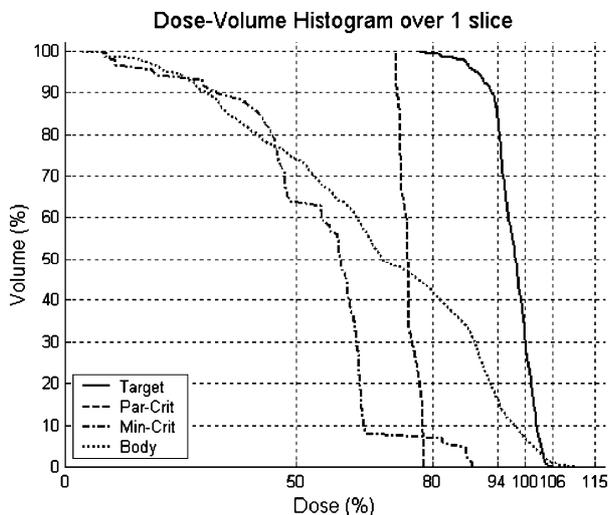


Fig. 9. DVH for IMRT, complex case.

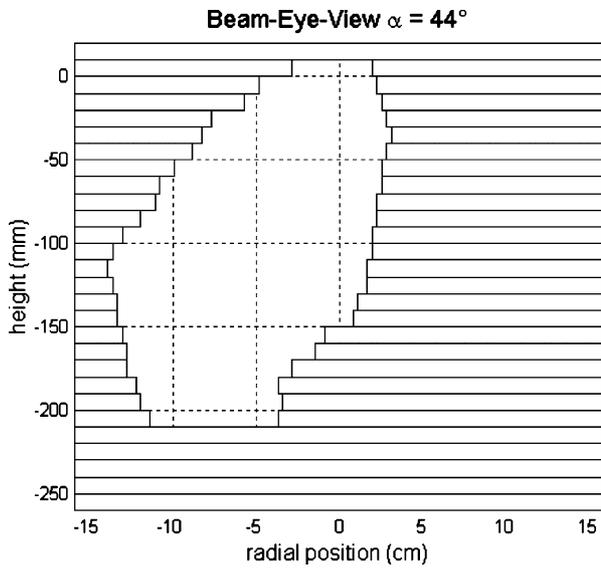


Fig. 11. Sample beam shape.

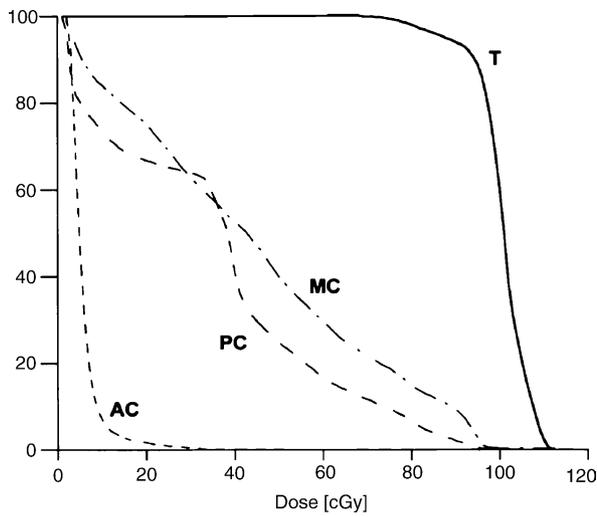


Fig. 12. DVHs calculated with commercial system.

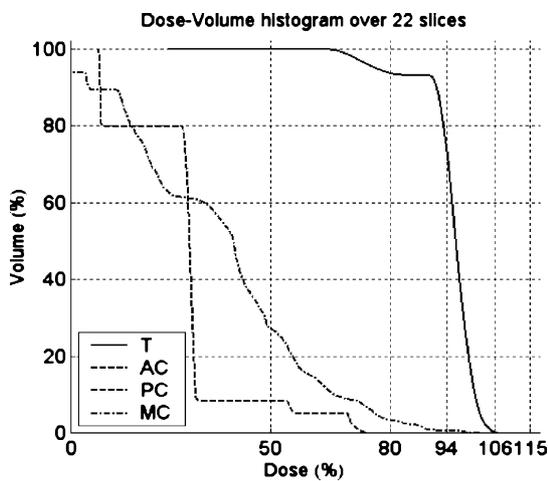


Fig. 13. DVHs calculated with designed software.

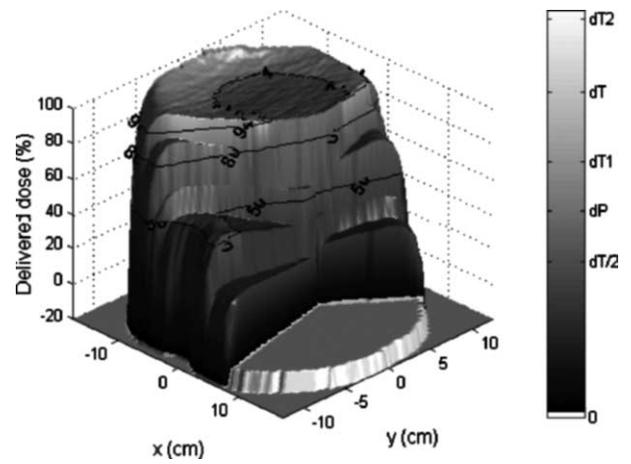


Fig. 14. 3D view of delivered dose over single slice.

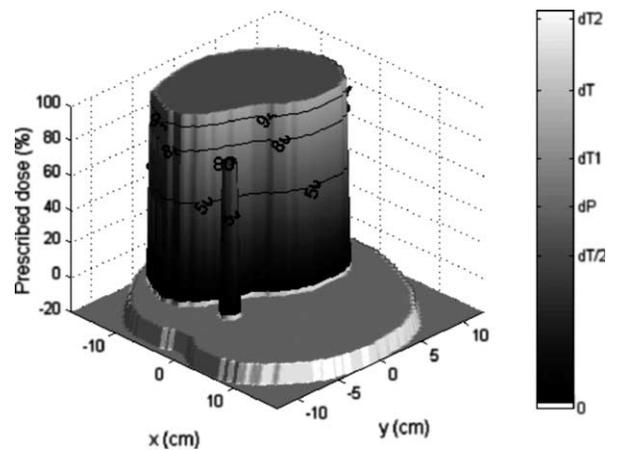


Fig. 15. 3D view of prescribed dose over single slice.

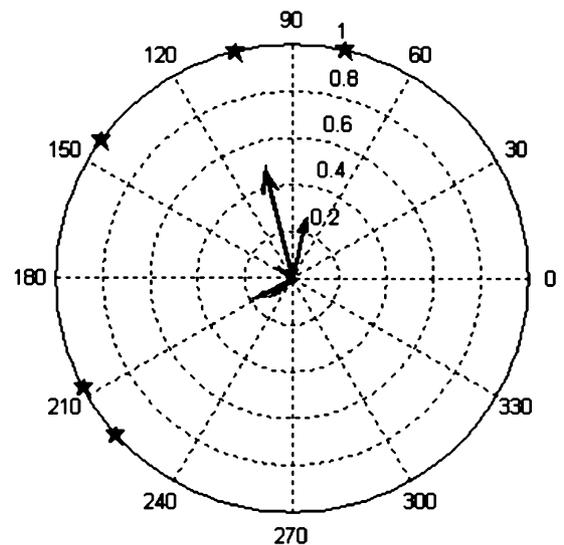


Fig. 16. Arrow-diagram representing beam configuration.

been validated through comparison with commercial system, and the whole optimisation algorithm has been tested with real clinical cases, both simple and complex. As stated by physician and by simulation with the same commercial system, our algorithm found good solutions in simple and complex cases using corresponding correct therapy, with direct chance to use them in real clinical treatment.

4. Conclusions

A unified evolutionary approach to inverse planning for three different kinds of coplanar radiotherapies has been proposed.

Thanks to genetic algorithm optimisation we have been able to search for full beam configurations, that is beam intensity, beam shape and especially beam orientation. To our best knowledge, only a few recent articles give a proposal solution for the angles optimisation and even fewer make use of evolutionary techniques.

Our implemented framework has the advantage of being simple and modular, with several small ‘parts’ that can be modified to take care of different situations encountered. Modelling 3D-CRT, ABRT and IMRT inverse planning with one common approach is a result of this.

References

- [1] S. Webb, Intensity-Modulated Radiation Therapy, Institute of Physics Publishing, Bristol, UK, 2000.
- [2] N. Milickovic, M. Lahanas, D. Baltas, N. Zamboglou, Intensity modulated beam radiation therapy dose optimisation with multiobjective evolutionary algorithms, Submitted for EMO.
- [3] G.A. Ezzell, Genetic and geometric optimisation of three-dimensional radiation therapy treatment planning, *Medical Physics* 23 (1996) 293–305.
- [4] J. Knowles, D. Corne, M. Bishop, Evolutionary training of artificial neural networks for radiotherapy treatment of cancers, *Proceedings of 1998 IEEE International Conference on Evolutionary Computation*, Alaska, 1998, pp. 398–403.
- [5] O.C.L. Haas, K.J. Burnham, J.A. Mills, Optimisation of beam orientation in radiotherapy using planar geometry, *Journal of Physics in Medicine and Biology* 43 (1998) 2179–2193.
- [6] E. Schreibmann, R. Uricchio, M. Lahanas, K. Theodorou, N. Zamboglou, C. Kappas, D. Baltas, A geometry based optimisation algorithm for conformal external beam orientation, Department of Medical Physics and Engineering, Klinikum Offenbach, Offenbach, Germany, 2001.
- [7] A. Pugachev, J.G. Li, A.L. Boyer, S.L. Hancock, Q.T. Le, S.S. Donaldson, L. Xing, Role of beam orientation optimisation in intensity-modulated radiation therapy, *International Journal of Radiation, Oncology, Biology, Physics* 50 (2) (2001) 551–560.
- [8] M. Braunstein, R.Y. Levine, *Optimum Beam Configurations in Tomographic Intensity Modulated Radiation Therapy*, Spectral Sciences, Inc., Burlington, MA, 2000.
- [9] S. Webb, Optimizing the planning of intensity-modulated radiotherapy, *Physics in Medicine and Biology* 39 (1994) 2229–2246.
- [10] S. Marzi, M. Mattia, P. Del Giudice, B. Caccia, M. Benassi, Optimisation of intensity modulated radiation therapy: assessing the complexity of the problem, *Annali dell Istituto Superiore di Sanita* 37 (2) (2001) 225–230.
- [11] D. Beasley, R. Bull, R.R. Martin, An overview of genetic algorithms: part 1, fundamentals, *University Computing* 15 (1993) 58–69.
- [12] D.E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison Wesley, Reading, MA, USA, 1989.
- [13] Z. Michalewicz, *Genetic algorithms + Data Structures = Evolution Programs*, Springer, Berlin, 1992.
- [14] C. Cotruz, L. Xing, Segment-based inverse planning using genetic algorithm, *Physics in Medicine and Biology* 48 (2003) 2987–2998.
- [15] A. Pugachev, L. Xing, Incorporating prior knowledge into beam orientation optimization in IMRT, *International Journal of Radiation Oncology, Biology, Physics* 54 (2002) 1565–1574.