

Simulated Annealing applied to IMRT Beam Angle Optimization: a computational study

Joana Dias^{1,2}, Humberto Rocha², Brígida Ferreira³, Maria do Carmo Lopes^{3,4}

¹Faculdade de Economia, Universidade de Coimbra, Portugal

²Inesc-Coimbra, Portugal

³I3N, Departamento de Física, Universidade de Aveiro, Portugal

⁴Serviço de Física Médica, IPOC-FG, EPE, Coimbra, Portugal

Abstract

Electing irradiation directions to use in IMRT treatments is one of the first decisions to make in treatment planning. Beam angle optimization (BAO) is a difficult problem to tackle from the mathematical optimization point of view. It is highly non-convex, and optimization approaches based on gradient-descent methods will probably get trapped in one of the many local minima. Simulated Annealing (SA) is a local search probabilistic procedure that is known to be able to deal with multimodal problems. SA for BAO was retrospectively applied to ten clinical examples of treated cases of head-and-neck tumors signalized as complex cases where proper target coverage and organ sparing proved difficult to obtain. The number of directions to use was considered fixed and equal to 5 or 7. It is shown that SA can lead to solutions that significantly improve organ sparing, even considering a reduced number of angles, without jeopardizing tumor coverage.

Keywords: beam angle optimization; simulated annealing; metaheuristics

Introduction

Beam angle optimization (BAO) for IMRT treatments is known to be a very challenging problem: it is highly non-convex, non-linear, and with many local minima. Many different approaches have been proposed to tackle BAO, from linear programming and gradient search [1], to neighborhood search and metaheuristics approaches (see, for instance, [2-10]). There are also commercial treatment planning systems (TPS) that already include some BAO functionalities (e.g. Eclipse from Varian or Oncentra from Elekta).

Simulated Annealing (SA) can be interpreted as a local search probabilistic procedure that will try to find a global minimum of a cost function even in the presence of many local minima [11]. SA is inspired by the thermal process for obtaining low energy states of a solid in a heat bath in condensed matter physics [12]. It is possible to prove the asymptotic convergence of SA to the set of globally optimal solutions under some conditions [12]. Detailed descriptions of SA can be found in [11-15].

1
2
3
4 SA has been applied before to radiation treatment planning. Webb [16, 17] applied SA to 3D-Conformal
5 treatment planning. Morrill *et al* [18] use constrained SA, where the assessment of a given solution is
6 done resorting to the objective function value but checking whether or not the solution satisfies a set of
7 predetermined constraints. They test different objective functions (maximizing the probability of
8 complication-free treatment and maximizing the dose to the isocenter, subject to some dose-volume
9 constraints). Rosen *et al* [19] compare four different SA algorithms for conformal therapy treatment
10 planning, concluding that Very Fast Simulated Reannealing [20] gave the best results. Djajaputra *et al*
11 [21] describe a fast dose calculation engine that uses an approximate dose kernel and is embedded in a SA
12 algorithm for BAO. Bertsimas *et al* [22] propose a hybrid heuristic where simulated annealing is
13 combined with gradient information for BAO.
14
15
16
17
18
19
20
21

22 In this paper we describe an application of SA for IMRT beam angle optimization. One of the most
23 important features of SA is the neighborhood structure it uses. In this work, we consider several different
24 neighborhood structures, and make several computational tests considering a set of ten head-and-neck
25 cancer patients using 5 or 7 angles. The results obtained allow us to conclude that the neighborhood
26 structure that seems to be most promising is the dynamically dimensioned neighborhood [23]. This is a
27 neighborhood with a random and dynamic structure that changes as the algorithm progresses. As far as
28 the authors know, it is the first time that a SA with a dynamically dimensioned neighborhood is
29 considered for BAO, and it is also the first time that the behavior of SA with different underlying
30 neighborhood structures for BAO is studied.
31
32
33
34
35
36

37 In the rest of this paper we describe in detail the SA algorithm and the neighborhood structures
38 considered. We describe the computational tests that were performed, presenting and discussing the
39 results.
40
41
42
43

44 **Methods**

45 SA is a local search procedure that begins with any admissible solution to the problem. This is called the
46 current solution. It then randomly selects a solution in the neighborhood of the current solution. If this
47 solution is better than the current one, it becomes the current solution and the process is repeated. If this
48 solution is worse than the current one, it can still become the current solution with a given probability.
49 This is one of the SA features that allows the algorithm to escape from local minima. By allowing the
50 deterioration of the objective function value at some points during the algorithm's execution, it is possible
51 to jump to interesting regions of the searchable surface.
52
53
54
55
56

57 The probability of accepting a solution that is worse than the current one is defined by a non-increasing
58 function usually known as *temperature*. This name is related to the physical thermal process that inspired
59
60
61
62
63
64
65

1
2
3
4 the development of SA. The *temperature* decreases as the algorithm progresses, meaning that the
5 probability of accepting a worse solution decreases with the increase in the number of iterations of the
6 algorithm.
7

8
9 To be able to implement the SA algorithm, it is necessary to define: a data structure to represent a given
10 solution to the problem; the neighborhood; a way of determining the quality of a given solution; a
11 temperature function.
12
13

14 15 16 *Representation of the solutions*

17 In this work we considered a very simple data structure to represent a solution: a vector θ with n
18 elements, where n is the total number of angles used in the IMRT treatment. Each element of this vector
19 (a variable) will belong to the interval $[0^\circ, 360^\circ[$, and for ease in the exposition it is considered that all
20 elements of θ are in ascending order. When dealing with the BAO problem, two different interpretations
21 of the problem can be considered. It can be interpreted as a combinatorial problem, where we want to find
22 the best combination of n angles out of the set of all possible combinations. This usually implies a
23 discretization of the interval $[0^\circ, 360^\circ[$. It can also be interpreted as a problem with continuous variables.
24 The approach followed in this paper is the latter. Despite the fact that each variable has to belong to
25 $[0^\circ, 360^\circ[$, it is not even necessary to consider an upper or lower bound for each variable, since an angle of
26 -10° , for instance, is equal to 350° or an angle of 370° is equal to 10° .
27
28
29
30
31
32
33
34
35

36 *Neighborhood Structures*

37 Two solutions are said to be *k-neighbors* if they have at most k different angles. Given a current solution,
38 a k -neighborhood solution is calculated by randomly generating at most k integer values belonging to
39 $[1, n]$, and then randomly generating values in the interval $[0^\circ, 360^\circ[$ that will be used to perturb the current
40 values of the chosen k variables. These perturbation values can be randomly generated using different
41 probability distributions: following a normal distribution with mean 0 and a given standard deviation r or
42 following a uniform distribution in $[0^\circ, 360^\circ[$, for instance. If using a normal distribution, it is also possible
43 to dynamically adjust the standard deviation.
44
45
46
47
48
49

50 The value of k can be defined *a priori* and be constant throughout the algorithm's execution, or it can be
51 dynamically and randomly determined. Both approaches were considered. In the latter case, we have
52 decided to test a dynamically dimensioned neighborhood [23, 24]. In such a neighborhood structure, the
53 value of k is randomly determined in each iteration. In reality, each angle in θ is or is not changed based
54 on a given probability. This probability is determined by a non-increasing function, that guarantees that
55 neighborhoods of greater dimension are considered in the beginning of the algorithm's execution,
56
57
58
59
60
61
62
63
64
65

decreasing the size of the neighborhoods as the algorithm progresses. In the present work, this probability function is given by the temperature function.

Assessment of the solutions

The assessment of a given solution to BAO can only be done after calculating the optimal fluence map for the set of angles represented by the solution. This means that this assessment requires the resolution of the fluence map optimization (FMO) problem, where the treatment angles are considered as an input.

There are many ways of optimizing IMRT fluence maps. It is out of the scope of this work to describe the pros and cons of the existing models and methods. We have chosen to use a convex penalty function voxel-based nonlinear model based on [25] where each voxel is penalized considering the square difference of the amount of dose received by the voxel and the amount of dose desired/allowed for the voxel. This formulation yields a programming problem with only linear nonnegativity constraints on the fluence values.

Considering that V equals the number of voxels, N equals the number of beamlets (considering only angles in θ) and D represents the dose matrix, such that D_{ij} represents the contribution of beamlet j to the total dose deposited in voxel i , we can say that the total dose received by voxel i is given by $\sum_{j=1}^N D_{ij} w_j$

with w_j representing the weight of beamlet j .

Let T_i be the desired dose for voxel i , $\underline{\lambda}_i$ and $\bar{\lambda}_i$ the penalty weights of underdose and overdose of voxel i , respectively, and $(\bullet)_+ = \max\{0, \bullet\}$. Then the FMO model can be defined as follows:

$$f(\theta) = \text{Min}_w \sum_{i=1}^V \left[\underline{\lambda}_i \left(T_i - \sum_{j=1}^N D_{ij} w_j \right)_+^2 + \bar{\lambda}_i \left(\sum_{j=1}^N D_{ij} w_j - T_i \right)_+^2 \right] \quad (1)$$

$$\text{s.t. } w_j \geq 0, j = 1, \dots, N \quad (2)$$

Although this formulation allows unique weights for each voxel, weights are assigned by structure only so that every voxel in a given structure has the weight assigned to that structure divided by the number of voxels of the structure. This nonlinear formulation implies that a very small amount of underdose or overdose may be accepted, but larger deviations from the desired/allowed doses are decreasingly tolerated.

For every solution considered, the FMO is performed, and the corresponding objective function value is used as a measure of the quality of the solution (the lower the better).

Temperature function

The temperature function, that will help determining the probability of accepting a worse neighboring solution at each iteration of the algorithm, can be any non-increasing function. We have chosen to work with a very simple logarithmic function of the form $T(i) = (1 - \log(i)/\log(N))$, where N represents the maximum number of iterations (Figure 1) and i represents the current iteration. The probability of accepting a worse solution than the current one is determined not only by the temperature function but also by the difference between the two objective function values. This probability is calculated as

$$e^{\frac{-(f(\theta_{new}) - f(\theta_{current}))}{T(i)f(\theta_{current})}}$$

. This means that solutions with an objective function value that is much worse than the current one will have a lower probability of being accepted as the current solution when compared with a solution that has an objective function value close to the current one.




Figure 1 Temperature function and Probability of accepting a worse solution

The SA algorithm

The SA algorithm will receive as input an initial solution $\theta_{current}$ and a temperature function T . It will return the best solution calculated (θ_{best}). The SA algorithm can be formally described as follows:

1. Set counter $i \leftarrow 1$. Define the initial solution $\theta_{current}$ and calculate $f(\theta_{current})$. $f_{best} \leftarrow f(\theta_{current})$; $\theta_{best} \leftarrow \theta_{current}$.
2. Generate randomly a new solution θ_{new} belonging to the neighborhood of $\theta_{current}$. Calculate $f(\theta_{new})$.
3. If $f(\theta_{new}) \leq f(\theta_{current})$, then $\theta_{current} \leftarrow \theta_{new}$. If $f(\theta_{new}) \leq f(\theta_{best})$, $f_{best} \leftarrow f(\theta_{new})$, $\theta_{best} \leftarrow \theta_{new}$.

4. If $f(\theta_{new}) > f(\theta_{current})$, then calculate $T(i)$ and $aux = e^{\frac{-(f(\theta_{new}) - f(\theta_{current}))}{T(i)f(\theta_{current})}}$. Generate randomly $p \in [0,1]$ using a uniform distribution. If $p \geq aux$ then $\theta_{current} \leftarrow \theta_{new}$.
5. $i \leftarrow i+1$. If the stopping criteria are met stop, else go to 2.

Step 1 initializes the iteration counter and the current solution. In Step 2, a solution is randomly selected from the neighborhood of the current solution. If this new generated solution is better than the current one, it will become the current solution. If it is also better than the best solution known so far, it will become the best solution and the best known value of the objective function is updated (Step 3). If this new solution is worse than the current solution, it can still be considered the current one in the next iteration with a given probability given by the temperature function (Step 4). The procedure is repeated until the stopping criteria are met (in the present case the stopping criterion was a maximum number of iterations).

A set of ten examples of clinical cases of head-and-neck tumors already treated at the Portuguese Institute of Oncology of Coimbra (IPOC) were used for computational tests. The selected clinical examples were signalized at IPOC as complex cases where proper target coverage and organ sparing, in particular parotid sparing (parotid glands are in close proximity to or even overlapping with the PTV), proved to be difficult to obtain. Figure 2 depicts the contoured structures for one CT slice of one representative patient. For simplicity, in this study, the OARs used for treatment optimization were limited to the spinal cord, the brainstem and the parotid glands. For the head-and-neck cases in study the PTV was separated in two parts with different prescribed doses: PTV1 and PTV2. The prescription dose for the target volumes and tolerance doses for the OARs considered in the optimization are presented in Table 1.

Table 1. Prescribed doses for all the structures considered for IMRT SA optimization

Structure	Mean dose	Maximum Dose	Prescribed Dose
Spinal cord	-	45 Gy	-
Brainstem	-	54 Gy	-
Left parotid	26 Gy	-	-
Right parotid	26 Gy	-	-
PTV1	-	-	70.0 Gy
PTV2	-	-	59.4 Gy
Body	-	80 Gy	-

Our tests were performed on a Intel Core i7 CPU 2.8 GHz computer with 4GB RAM and Windows 7. We used CERR 3.2.2 version ([26]) and MATLAB 7.4.0 (R2007a). The dose was computed using CERR's pencil beam algorithm (QIB), with corrections for heterogeneities. For each of the ten head-and-neck cases, the sample rate used for Body was 32 and for the remaining structures was 4 (meaning that each set of 32 Body voxels was considered as one voxel in the optimization procedure and, for all other structures,

one out of 4 voxels was used in the optimization procedure). To address the convex nonlinear formulation of the FMO problem we used a trust-region-reflective algorithm (*fmincon*) of MATLAB 7.4.0 (R2007a) Optimization Toolbox. Weights associated with the different structures were defined *a priori*, are the same for all patients, and were considered fixed during the algorithm's execution.



Figure 2 Contoured structures in one CT slice for one representative patient

In a first step, several different versions of the SA algorithm, considering different neighborhood structures, were tested considering IMRT treatments with 5 gantry angles. Simple local search procedures using the same neighborhood structures as the SA algorithm were also tested and compared with the SA approach, in a total of 12 different search procedures. After analyzing the results, the local search procedure that was considered the best one for the 5 angles BAO problem was then applied to the same set of patients but considering IMRT treatments with 7 angles.

Different local search procedures, 12 in total, were tested (Table 2).

Table 2. Local search procedures tested

Version of the algorithm	Description
1	SA with k -neighborhoods defined by k constant and equal to 1. Angles are changed by considering a perturbation that is randomly generated following a normal distribution with mean 0 and standard deviation r .
2	SA with k -neighborhoods defined by k constant and equal to 2. Angles are changed by considering a perturbation that is

	randomly generated following a normal distribution with mean 0 and standard deviation r .
3	SA with k -neighborhoods defined by k constant and equal to 5. Angles are changed by considering a perturbation that is randomly generated following a normal distribution with mean 0 and standard deviation r .
4	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 1. Angles are changed by considering a perturbation that is randomly generated following a normal distribution with mean 0 and standard deviation r .
5	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 2. Angles are changed by considering a perturbation that is randomly generated following a normal distribution with mean 0 and standard deviation r .
6	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 5. Angles are changed by considering a perturbation that is randomly generated following a normal distribution with mean 0 and standard deviation r .
7	SA with k -neighborhoods defined by k constant and equal to 1. Angles are changed by considering a perturbation that is randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ]$.
8	SA with k -neighborhoods defined by k constant and equal to 2. Angles are changed by considering a perturbation that is randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ]$.
9	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 1. Angles are changed by considering a perturbation that is randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ]$.
10	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 2. Angles are changed by considering a perturbation that is randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ]$.
11	SA with dynamically dimensioned neighborhood.
12	Dynamically dimensioned search without SA.

Considering previous computational experiments [24, 27], whenever using a normal distribution, we have decided to dynamically change the standard deviation r as follows:

1. If there are $l_{success}$ successful successive iterations (iterations where the best objective function is improved) then $r \leftarrow \min\{2r, rmax\}$.
2. If there are $l_{failure}$ unsuccessful successive iterations (iterations where the best objective function is not improved) then $r \leftarrow \max\{r/2, rmin\}$.

The standard deviation is initialized as $r \leftarrow 360^\circ / (4 * n)$. $rmax$ is considered equal to 90 and $rmin$ is considered equal to 3.

Considering the random behavior of these procedures, each one was executed five times for each patient. A maximum of 200 iterations were considered. The obtained solutions were compared with the 5 angles equidistant solution used in clinical practice.

Results

BAO considering 5 angles

The average improvement in the objective function value is depicted in Table 3, where the best result for each patient is highlighted.

The interpretation of these results is not straightforward. As a matter of fact, it is not possible to execute all optimization procedures for each patient and then choose the best solution found. This would take a prohibitive computational time. So, how should we choose what algorithm to use?

Table 3. Improvement in the objective function value for the different algorithms tested (Mean values). Shadow values show the best results for each patient.

patient	1	2	3	4	5	6	7	8	9	10	11	12
1	4,65%	4,81%	4,45%	4,26%	4,04%	5,07%	3,87%	4,06%	4,08%	2,96%	4,22%	4,37%
2	5,31%	6,00%	5,59%	5,10%	5,14%	4,22%	5,92%	4,24%	5,78%	4,45%	6,44%	5,15%
3	3,48%	3,56%	4,08%	4,23%	3,94%	3,85%	3,78%	3,07%	3,06%	3,16%	3,75%	3,43%
4	5,46%	5,14%	5,26%	5,23%	5,51%	4,88%	4,03%	3,38%	4,50%	4,00%	4,96%	4,87%
5	10,23%	9,92%	10,07%	10,05%	9,80%	9,25%	9,17%	9,01%	10,61%	9,37%	10,46%	9,77%
6	6,07%	6,25%	5,81%	6,63%	6,78%	5,62%	5,93%	5,88%	5,87%	5,95%	5,78%	6,24%
7	9,89%	7,99%	8,80%	11,01%	10,96%	10,10%	10,35%	9,30%	8,25%	10,39%	12,09%	9,92%
8	9,09%	6,99%	8,77%	9,30%	8,08%	9,21%	7,26%	6,86%	8,48%	9,78%	9,14%	8,44%
9	9,52%	8,06%	9,06%	9,02%	8,49%	9,08%	8,13%	7,22%	9,61%	10,12%	9,10%	8,06%
10	1,96%	1,34%	1,48%	1,66%	1,39%	1,23%	0,84%	0,53%	0,93%	0,20%	3,04%	2,19%
Average improvement	6,56%	6,01%	6,33%	6,65%	6,41%	6,25%	5,93%	5,35%	6,12%	6,04%	6,90%	6,24%

There are several different ways of choosing a given algorithm from a set of possible algorithms. One possibility is by looking at the average improvements and choosing the algorithm that presents the best average behavior over the set of 10 patients. Applying this criterion, SA with dynamically dimensioned search would be the best choice. The problem of choosing an algorithm looking at average values is that a very bad result for a given patient could be compensated by medium quality results in other patients. So, another possibility could be to iteratively eliminate the algorithm that leads to the worst result for most patients (and choosing the one with the worst average value in case of a tie). Following this procedure, we would first eliminate procedure 8, since it leads to the worst result for 4 out of the 10 patients. We would then eliminate procedure 10, then 7, and so on. We would end up by also choosing procedure 11.

Still another possibility would be to use a kind of cross-validation procedure to select the best algorithm for a given patient. We would partition the set of patients such that one patient at a time is considered the “new” patient for which an algorithm will have to be chosen. The algorithm chosen would be the one with the best average results for the remaining 9 patients. If this procedure is applied here, we would choose algorithm 11 for all patients.

Although the SA with dynamically dimensioned search is not the best algorithm for all patients, it seems to be the one that would most probably be chosen to be applied to a new patient since it is the one chosen under different criteria.

Although we are assessing each solution using the optimal objective function of the FMO problem, more important than the objective function value is the dosimetric outcome of the generated plans. Table 4 and Table 5 summarize the dosimetric results. A metric usually used for plan evaluation is the dose in 95% of the PTV’s volume. Typically, we would like that 95% of the volume receives at least 95% of the

prescribed dose. In order to verify parotids' sparing, mean doses are usually considered. For spinal cord and brainstem, maximum doses cannot be exceeded. These metrics are displayed for the ten cases, considering the equidistant solution and the best and worst solutions out of the 5 solutions generated for each patient. The choice of the *best* and *worst* solutions is done resorting to the best and worst value of the corresponding FMO objective function. As this function is a weighted sum of deviations, there is no guarantee that the *best* solution is indeed better than the *worst* solution considering the dose metrics. The SA generated solutions do not present huge differences from the equidistant solution. However, we can see a better tumor dosimetry, and a better sparing especially of the left parotids. This is sometimes achieved by higher levels of maximum dose to the spinal cord and brainstem, but the achieved values are still within the tolerance doses for all but one patient.

Table 4. Comparison of dose-volume metrics using SA with dynamically dimensioned neighborhoods and 5 equidistant angles treatment plans

patient	$D_{95\%} (PTV_{70})/Gy$			$D_{95\%} (PTV_{59.4})/Gy$		
	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst
1	65,7	65,9	65,9	57,5	57,3	57,3
2	67,0	67,1	67,1	55,8	56,3	56,2
3	66,3	66,5	66,5	56,7	56,9	56,7
4	65,4	65,2	65,4	55,3	55,7	55,9
5	66,7	67,0	67,1	55,7	56,2	55,9
6	65,8	66,0	65,9	56,9	57,4	57,4
7	68,0	68,2	68,1	57,5	57,8	57,8
8	67,4	67,7	67,7	56,5	56,9	56,9
9	67,2	67,0	67,0	57,2	57,5	57,5
10	66,0	66,3	66,1	56,9	57,0	56,9

Table 5. Comparison of OARs dose metrics using SA with dynamically dimensioned neighborhoods and 5 equidistant angles treatment plans

patient	$D_{mean} (Right Parotid)/Gy$			$D_{mean} (Left Parotid)/Gy$			$D_{max} (Spinal Cord)/Gy$			$D_{max} (Brainstem)/Gy$		
	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst
1	25,6	23,1	23,1	26,4	25,1	24,7	38,9	40,6	40,9	52,6	52,3	51,9
2	25,2	26,7	24,4	26,4	26,2	25,6	45,0	44,9	40,8	55,2	55,0	54,8
3	26,2	25,4	24,8	27,3	27,4	26,0	42,5	39,6	43,5	40,6	45,4	46,6
4	28,2	29,0	28,4	27,1	26,1	26,2	40,1	42,5	42,3	50,0	49,2	50,2
5	29,1	30,7	29,5	26,5	23,3	25,0	38,5	39,1	39,9	50,3	51,6	51,9
6	25,7	22,2	22,4	24,9	22,7	23,5	38,7	40,7	40,3	52,0	52,9	53,0
7	26,4	27,4	26,6	26,4	26,9	27,5	41,0	41,9	40,5	52,3	51,6	51,3
8	27,1	23,2	27,2	25,3	24,7	24,2	39,7	40,3	38,4	52,4	52,6	54,0
9	20,8	21,1	20,0	25,2	25,1	24,2	39,4	40,9	42,7	51,4	50,4	50,7
10	26,4	25,1	26,2	24,9	25,4	25,2	41,1	39,7	40,4	50,4	52,8	49,8

1
2
3
4 *BAO considering 7 angles*

5
6 The algorithm considered as the best one for the 5 angle BAO problem (SA with dynamically
7 dimensioned neighborhoods) was applied to the same set of patients, but considering 7 beams. The
8 algorithm was executed 5 times for each patient, and the results were then compared with the 7 angles
9 equidistant solution. Figure 3 presents a boxplot comparing the results obtained for the PTVs in the 5
10 executions of the SA algorithm and the equidistant solution. The horizontal lines represent 95% of the
11 prescribed dose. Satisfactory treatment plans should obtain results above these lines. In average, the
12 optimized SA solutions present slightly better tumor coverage. Figure 4 presents similar charts
13 considering dosimetric results for the parotids. The horizontal lines represent the desirable maximum
14 mean dose for the corresponding structure. Satisfactory treatment plans should obtain results below these
15 lines. This is a difficult goal to achieve and, in this study, it is was only possible to obtain treatments plans
16 complying with this objective in about half the patients. It was possible to improve parotids' sparing in 6
17 out of 10 patients. For all generated plans it was possible to comply with the medical prescription for the
18 brainstem and spinal cord organs in 9 out of 10 patients (Figure 5).

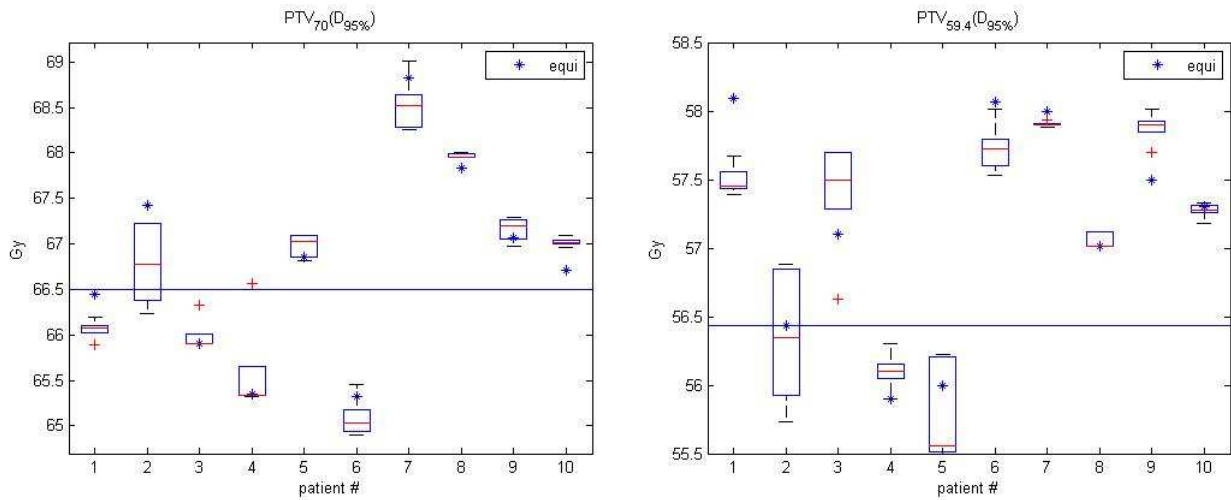


Figure 3 Comparison of dose-volume metrics using SA with dynamically dimensioned neighborhoods and 7 equidistant angles treatment plans

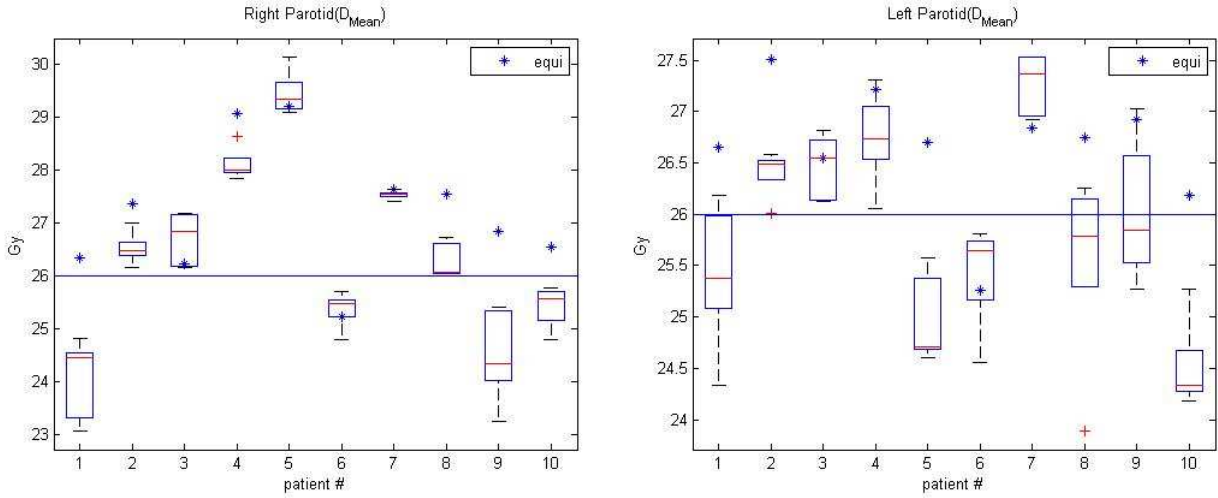


Figure 4 Comparison of parotids' dosimetry metrics using SA with dynamically dimensioned neighborhoods and 7 equidistant angles treatment plans

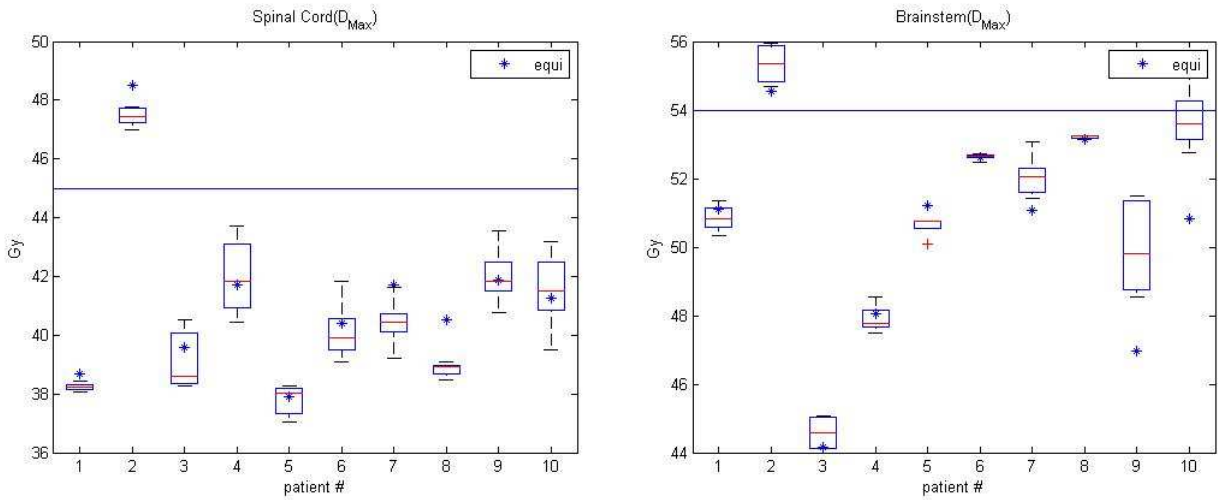


Figure 5 Comparison of spinal cord and brainstem dose using SA with dynamically dimensioned neighborhoods and 7 equidistant angles treatment plans

For patient 10, the parotids dose volume histograms of one of the SA solutions are compared with the 7 equidistant angles solution (Figure 6). Dose volume histograms are also shown for PTVs (Figure 7), spinal cord and brainstem (Figure 8). As can be seen, the irradiation of the PTVs achieved by the SA solution and the equidistant solution are almost the same.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 6 Dose Volume Histogram for Parotids




Figure 7 Dose Volume Histogram for PTVs




Figure 8 Dose Volume Histogram for Spinal Cord and Brainstem

Regarding computational times needed to run the SA algorithm, most of the computational time is spent solving the FMO problem for each solution. So, the computational time does not present significant differences for different neighborhood structures. For this set of patients and with the described computational resources, each instance of the FMO problem can take from 56 seconds to 350 seconds to be calculated, depending on the patient and on the set of beam angles considered. This corresponds to a mean total computational time of approximately 6 hours, and a maximum computational time of 11 hours. This computational time would be significantly decreased with the use of more powerful computational resources and parallel programming. It should be noticed that the SA procedure does not need any type of human intervention, so it can be run during night periods, or when the planner is occupied with other tasks.

Discussion

The computational results show that it is possible to improve organ sparing without jeopardizing tumor coverage by applying BAO. SA was able to generate solutions that lead to better plans than the usual equidistant solution. SA has as main advantages the fact that it is based on a simple algorithm, easy to translate into computational programming, it does not require many computational resources (like memory), and it is very flexible in the sense that it is possible to consider very different ways of assessing

1
2
3
4 the quality of a solution. Its main drawback, especially important for the IMRT BAO problem, is the
5 computational time needed if the objective function considered is computationally expensive (as is the
6 case with the FMO). This drawback made us consider a maximum number of iterations as the termination
7 criterion. Further improvements can be obtained by parallelizing the SA implementation. This will make
8 it possible to visit more than one neighboring solution in each iteration, or to apply a multi-start search,
9 where several SA algorithms are run in parallel, each one having a different initial solution. This will
10 contribute to a significant improvement in the obtained results.

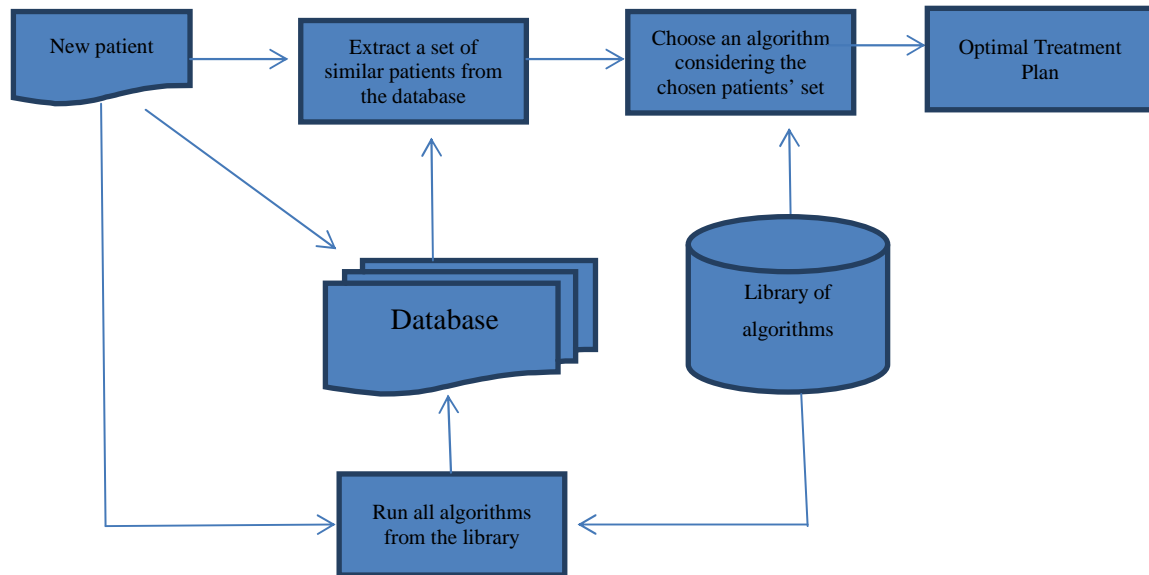
11
12
13
14
15
16 The random behavior of the SA algorithm can cause some concerns if applied to clinical practice. It is not
17 acceptable that a plan has unexpected changes after re-optimization because of a random behavior of the
18 algorithm. Actually, computers can only generate pseudo-random numbers. This makes it possible to
19 guarantee the calculation of exactly the same treatment plan every time the algorithm is executed, being
20 only necessary to consider a constant seed when initializing the random numbers generator. We did not
21 choose to do that in this paper, because it could lead to biased results: showing results for several runs is
22 the correct thing to do in this context.

23
24
25
26
27 In this paper, we decided to fix the weights used in the FMO *a priori*. Furthermore, these weights are the
28 same for all patients considered. It would be possible to improve considerably the results of the SA
29 algorithm if the choice of weights was patient dependent. Considering the particular features of each
30 patient, and after some trial-and-error attempts, it would be possible to define weights that would
31 guarantee better organ sparing and/or better target irradiation. Nevertheless, as the focus of this paper is to
32 compare different neighborhood structures for SA, and compare SA solutions with the equidistant
33 solution, working with different weights for different patients could introduce bias in the results, and
34 would certainly favor the SA approach over the equidistant solution. We are presently working in
35 automated ways for determining the weights to use in FMO.

36
37
38
39
40
41 Comparing the different solutions obtained makes it clear that the assessment of a given solution is
42 difficult to make looking only at a single objective function value. Further developments of this work will
43 consider incorporating a multiobjective approach in SA optimization.

44
45
46
47
48
49 As it is not possible to elect one single optimization algorithm as the best one for all patients, and is also
50 not possible to run several different optimization algorithms for every incoming patient due to time
51 constraints, a selection procedure based on cross-validation could be an interesting approach. In such an
52 approach, the optimization algorithm to be applied would be chosen considering the one that gave the best
53 average results in a set of already known patients. This set of already known patients could be chosen by
54 applying clustering techniques that would select the set of patients most similar to the incoming patient.
55 This would require the creation of a library of optimization algorithms, that would be applied to all known
56
57
58
59
60
61
62
63
64
65

1
2
3
4 patients so that all resulting data could be saved in a database. It should be noted that there is not enough
5 time to run several optimization algorithms for a new patient that has to have a treatment plan as soon as
6 possible. But it is possible to run several optimization algorithms, only for database recording purposes,
7 for each patient in the database since time is no longer important in this context (Figure 9).
8
9
10
11
12
13



34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Figure 9 Selection of optimization algorithm based on a cross-validation procedure

Conclusions

In this paper several different neighborhoods to be used in a SA algorithm applied to BAO were tested. It was possible to conclude that the best neighborhood to use was the dynamically dimensioned neighborhood, that is characterized by considering larger neighborhoods in the beginning of the algorithm's execution, progressing towards smaller neighborhoods. It is also possible to conclude that the major advantages of using SA applied to BAO are related more to organs' sparing than to a better irradiation of the PTVs.

Acknowledgments

This work was supported by FEDER, COMPETE, iCIS (CENTRO-07-ST24-FEDER-002003), Portuguese Foundation for Science and Technology under project grants PEst-OE/EEI/UI308/2014FCT, PTDC/EIA-CCO/121450/2010.

References

- [1] Craft, D.: Local beam angle optimization with linear programming and gradient search. *Physics in Medicine and Biology*. 52, N127-N135 (2007)
- [2] Li, Y., Yao, J., Yao, D.: Automatic beam angle selection in IMRT planning using genetic algorithm. *Physics in Medicine and Biology*. 49, 1915 (2004)
- [3] Li, Y., Yao, D., Yao, J., Chen, W.: A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning. *Physics in Medicine and Biology*. 50, 3491 (2005)
- [4] Lim, G.J., Cao, W.: A two-phase method for selecting IMRT treatment beam angles: Branch-and-Prune and local neighborhood search. *European Journal of Operational Research*. 217, 609-618 (2012)
- [5] Lei, J., Li, Y. *A DNA genetic algorithm for beam angle selection in radiotherapy planning*. in *Cybernetics and Intelligent Systems, 2008 IEEE Conference on*. 2008: IEEE.
- [6] Nazareth, D.P., Brunner, S., Jones, M.D., Malhotra, H.K., Bakhtiari, M.: Optimization of beam angles for intensity modulated radiation therapy treatment planning using genetic algorithm on a distributed computing platform. *Journal of Medical Physics/Association of Medical Physicists of India*. 34, 129 (2009)
- [7] Li, Y., Yao, D., Zheng, J., Yao, J. *A modified genetic algorithm for the beam angle optimization problem in intensity-modulated radiotherapy planning*. in *Artificial Evolution*. 2006: Springer.
- [8] Zhang, H.H., Gao, S., Chen, W., Shi, L., D'Souza, W.D., Meyer, R.R.: A surrogate-based metaheuristic global search method for beam angle selection in radiation treatment planning. *Physics in Medicine and Biology*. 58, 1933-1946 (2013)
- [9] Dias, J., Rocha, H., Ferreira, B.C., Lopes, M.C.: A genetic algorithm with neural network fitness function evaluation for IMRT beam angle optimization. *Central European Journal of Operations Research*. 22, 431-455 (2014)
- [10] Rocha, H., Dias, J., Ferreira, B.C., Lopes, M.C.: Selection of intensity modulated radiation therapy treatment beam directions using radial basis functions within a pattern search methods framework. *Journal of Global Optimization*. 57, 1065-1089 (2013)
- [11] Bertsimas, D., Tsitsiklis, J.: Simulated annealing. *Statistical Science*. 10-15 (1993)
- [12] Aarts, E., Korst, J., Michiels, W.: Simulated Annealing. In: (eds) *Search Methodologies*. vol. pp. 188-210. Springer, (2005)
- [13] Kirkpatrick, S., Jr, D.G., Vecchi, M.P.: Optimization by simulated annealing. *Science*. 220, 671-680 (1983)
- [14] Ingber, L.: Simulated annealing: Practice versus theory. *Mathematical and Computer Modelling*. 18, 29-57 (1993)
- [15] Webb, S.: Optimizing radiation therapy inverse treatment planning using the simulated annealing technique. *International Journal of Imaging Systems and Technology*. 6, 71-79 (1995)
- [16] Webb, S.: Optimization by simulated annealing of three-dimensional, conformal treatment planning for radiation fields defined by a multileaf collimator: II. Inclusion of two-dimensional modulation of the x-ray intensity. *Physics in Medicine and Biology*. 37, 1689 (1992)
- [17] Webb, S.: Optimization by simulated annealing of three-dimensional conformal treatment planning for radiation fields defined by a multileaf collimator. *Physics in Medicine and Biology*. 36, 1201-1226 (1991)
- [18] Morrill, S.M., Lane, R.G., Jacobson, G., Rosen, II: Treatment planning optimization using constrained simulated annealing. *Physics in Medicine and Biology*. 36, 1341-1361 (1991)
- [19] Rosen, I.I., Lam, K.S., Lane, R.G., Langer, M., Morrill, S.M.: Comparison of simulated annealing algorithms for conformal therapy treatment planning. *International Journal of Radiation Oncology, Biology, Physics*. 33, 1091-1099 (1995)
- [20] Morrill, S.M., Lam, K.S., Lane, R.G., Langer, M., Rosen, I.I.: Very fast simulated reannealing in radiation therapy treatment plan optimization. *International Journal of Radiation Oncology, Biology, Physics*. 31, 179-188 (1995)

- 1
2
3
4 [21] Djajaputra, D., Wu, Q., Wu, Y.,Mohan, R.: Algorithm and performance of a clinical IMRT beam-
5 angle optimization system. *Physics in Medicine and Biology*. 48, 3191-3212 (2003)
6 [22] Bertsimas, D., V.Cacchiani, D.Craft,O.Nohadani: A hybrid approach to beam angle optimization in
7 intensity-modulated radiation therapy. *Computers & Operations Research*. 40, 2187-2197 (2013)
8 [23] Tolson, B.,Shoemaker, C.: Dynamically dimensioned search algorithm for computationally efficient
9 watershed model calibration. *Water Resources Research*. 43, W01413 (2007)
10 [24] Regis, R.G.,Shoemaker, C.A.: Combining radial basis function surrogates and dynamic coordinate
11 search in high-dimensional expensive black-box optimization. *Engineering Optimization*. 45, 529-555
12 (2013)
13 [25] Aleman, D.M., Kumar, A., Ahuja, R.K., Romeijn, H.E.,Dempsey, J.F.: Neighborhood search
14 approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning.
15 *Journal of Global Optimization*. 42, 587-607 (2008)
16 [26] Deasy, J.O., Blanco, A.I.,Clark, V.H.: CERR: A computational environment for radiotherapy
17 research. *Medical Physics*. 30, 979-985 (2003)
18 [27] Dias, J., Rocha, H., Ferreira, B.,Lopes, M.C.: IMRT Beam Angle Optimization Using DDS with a
19 Cross-Validation Approach for Configuration Selection. In: B. Murgante, S. Misra, A.M.A.C. Rocha, C.
20 Torre, J.G. Rocha, M.I. Falcão, D. Taniar, B.O. Apduhan, O. Gervasi (eds) *Computational Science and Its*
21 *Applications – ICCSA 2014*. Lecture Notes in Computer Science, vol. 8520, pp. 1-16. Springer, (2014)
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Table 1. Prescribed doses for all the structures considered for IMRT SA optimization

Structure	Mean dose	Maximum Dose	Prescribed Dose
Spinal cord	–	45 Gy	–
Brainstem	–	54 Gy	–
Left parotid	26 Gy	–	–
Right parotid	26 Gy	–	–
PTV1	–	–	70.0 Gy
PTV2	–	–	59.4 Gy
Body	–	80 Gy	–

Table 2. Local search procedures tested

Version of the algorithm	Description
1	SA with k -neighborhoods defined by k constant and equal to 1. Angles are randomly generated following a normal distribution with mean 0 and standard deviation r .
2	SA with k -neighborhoods defined by k constant and equal to 2. Angles are randomly generated following a normal distribution with mean 0 and standard deviation r .
3	SA with k -neighborhoods defined by k constant and equal to 5. Angles are randomly generated following a normal distribution with mean 0 and standard deviation r .
4	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 1. Angles are randomly generated following a normal distribution with mean 0 and standard deviation r .
5	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 2. Angles are randomly generated following a normal distribution with mean 0 and standard deviation r .
6	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 5. Angles are randomly generated following a normal distribution with mean 0 and standard deviation r .
7	SA with k -neighborhoods defined by k constant and equal to 1. Angles are randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ[$.
8	SA with k -neighborhoods defined by k constant and equal to 2. Angles are randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ[$.
9	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 1. Angles are randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ[$.
10	Local search procedure without SA. k -neighborhoods defined by k constant and equal to 2. Angles are randomly generated following a uniform distribution in the interval $[0^\circ, 360^\circ[$.
11	SA with dynamically dimensioned neighborhood.
12	Dynamically dimensioned search without SA.

Table 3. Improvement in the objective function value (Mean values)

patient	1	2	3	4	5	6	7	8	9	10	11	12
1	4,65%	4,81%	4,45%	4,26%	4,04%	5,07%	3,87%	4,06%	4,08%	2,96%	4,22%	4,37%
2	5,31%	6,00%	5,59%	5,10%	5,14%	4,22%	5,92%	4,24%	5,78%	4,45%	6,44%	5,15%
3	3,48%	3,56%	4,08%	4,23%	3,94%	3,85%	3,78%	3,07%	3,06%	3,16%	3,75%	3,43%
4	5,46%	5,14%	5,26%	5,23%	5,51%	4,88%	4,03%	3,38%	4,50%	4,00%	4,96%	4,87%
5	10,23%	9,92%	10,07%	10,05%	9,80%	9,25%	9,17%	9,01%	10,61%	9,37%	10,46%	9,77%
6	6,07%	6,25%	5,81%	6,63%	6,78%	5,62%	5,93%	5,88%	5,87%	5,95%	5,78%	6,24%
7	9,89%	7,99%	8,80%	11,01%	10,96%	10,10%	10,35%	9,30%	8,25%	10,39%	12,09%	9,92%
8	9,09%	6,99%	8,77%	9,30%	8,08%	9,21%	7,26%	6,86%	8,48%	9,78%	9,14%	8,44%
9	9,52%	8,06%	9,06%	9,02%	8,49%	9,08%	8,13%	7,22%	9,61%	10,12%	9,10%	8,06%
10	1,96%	1,34%	1,48%	1,66%	1,39%	1,23%	0,84%	0,53%	0,93%	0,20%	3,04%	2,19%
Average improvement	6,56%	6,01%	6,33%	6,65%	6,41%	6,25%	5,93%	5,35%	6,12%	6,04%	6,90%	6,24%

Table 4. Comparison of dose-volume metrics using SA with dynamically dimensioned neighborhoods and 5 equidistant angles treatment plans

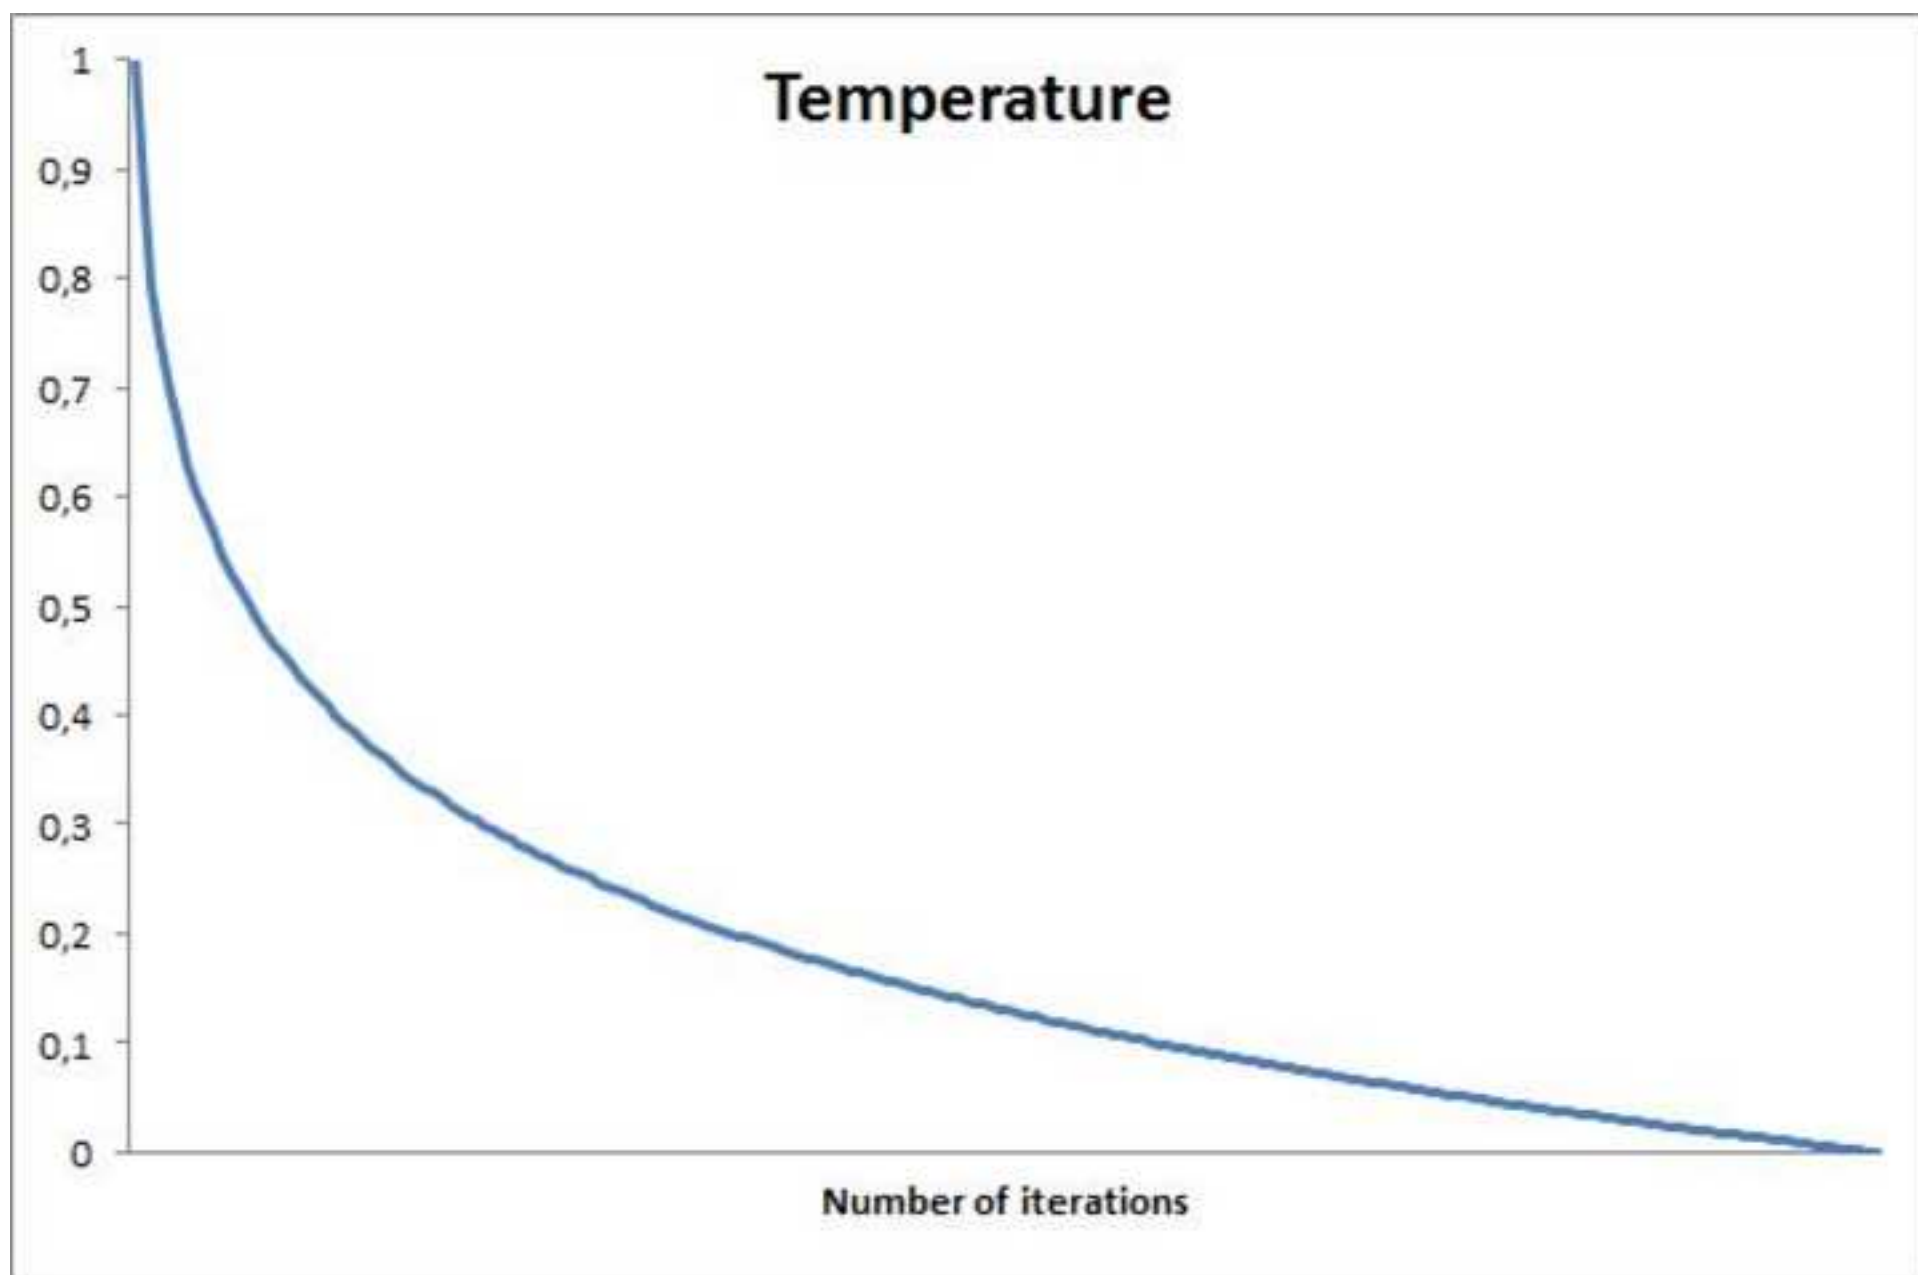
patient	$D_{95\%} (PTV_{70})/Gy$			$D_{95\%} (PTV_{59.4})/Gy$		
	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst
1	65,7	65,9	65,9	57,5	57,3	57,3
2	67,0	67,1	67,1	55,8	56,3	56,2
3	66,3	66,5	66,5	56,7	56,9	56,7
4	65,4	65,2	65,4	55,3	55,7	55,9
5	66,7	67,0	67,1	55,7	56,2	55,9
6	65,8	66,0	65,9	56,9	57,4	57,4
7	68,0	68,2	68,1	57,5	57,8	57,8
8	67,4	67,7	67,7	56,5	56,9	56,9
9	67,2	67,0	67,0	57,2	57,5	57,5
10	66,0	66,3	66,1	56,9	57,0	56,9

Table 5. Comparison of OARs dose metrics using SA with dynamically dimensioned neighborhoods and 5 equidistant angles treatment plans

patient	D _{mean} (Righ Parotid)/Gy			D _{mean} (Left Parotid)/Gy			D _{max} (Spinal Cord)/Gy			D _{max} (Brainstem)/Gy		
	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst	Equi	SA_best	SA_worst
1	25,6	23,1	23,1	26,4	25,1	24,7	38,9	40,6	40,9	52,6	52,3	51,9
2	25,2	26,7	24,4	26,4	26,2	25,6	45,0	44,9	40,8	55,2	55,0	54,8
3	26,2	25,4	24,8	27,3	27,4	26,0	42,5	39,6	43,5	40,6	45,4	46,6
4	28,2	29,0	28,4	27,1	26,1	26,2	40,1	42,5	42,3	50,0	49,2	50,2
5	29,1	30,7	29,5	26,5	23,3	25,0	38,5	39,1	39,9	50,3	51,6	51,9
6	25,7	22,2	22,4	24,9	22,7	23,5	38,7	40,7	40,3	52,0	52,9	53,0
7	26,4	27,4	26,6	26,4	26,9	27,5	41,0	41,9	40,5	52,3	51,6	51,3
8	27,1	23,2	27,2	25,3	24,7	24,2	39,7	40,3	38,4	52,4	52,6	54,0
9	20,8	21,1	20,0	25,2	25,1	24,2	39,4	40,9	42,7	51,4	50,4	50,7
10	26,4	25,1	26,2	24,9	25,4	25,2	41,1	39,7	40,4	50,4	52,8	49,8


Figure

[Click here to download high resolution image](#)



Figure

[Click here to download high resolution image](#)



Figure

[Click here to download high resolution image](#)



Figure
[Click here to download high resolution image](#)




Figure
[Click here to download high resolution image](#)




Figure
[Click here to download high resolution image](#)




Figure
[Click here to download high resolution image](#)




Figure
[Click here to download high resolution image](#)

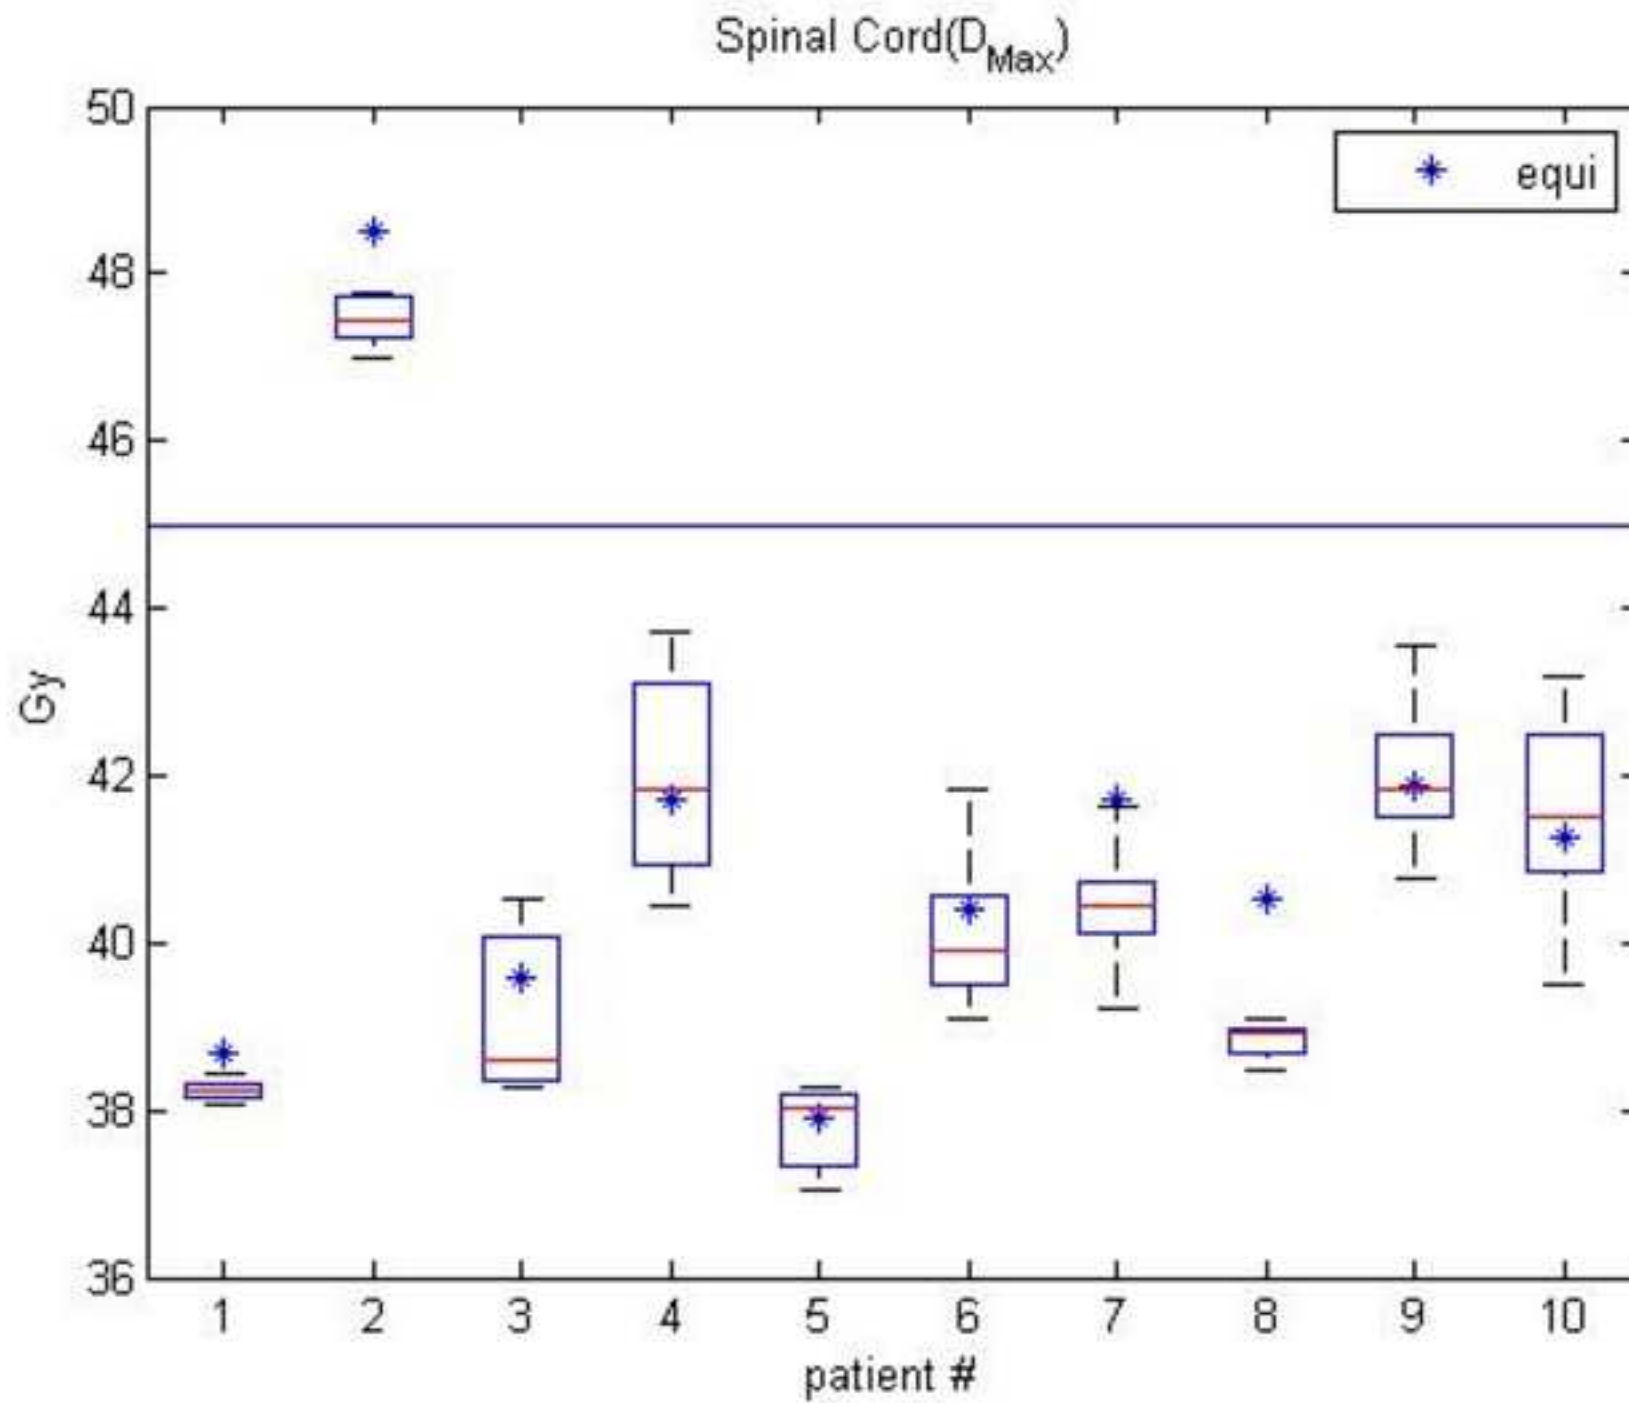




Figure
[Click here to download high resolution image](#)




Figure

[Click here to download high resolution image](#)



Figure

[Click here to download high resolution image](#)



Figure

[Click here to download high resolution image](#)




Figure
[Click here to download high resolution image](#)

