

IMRT Beam Angle Optimization Using DDS with a Cross-Validation Approach for Configuration Selection

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

¹Faculdade de Economia, Universidade de Coimbra, 3004-512 Coimbra, Portugal

²Inesc-Coimbra, Rua Antero de Quental, 199, 3000-033 Coimbra, Portugal

³IN, Departamento de Física, Universidade de Aveiro, 3810-193 Aveiro, Portugal

⁴Serviço de Física Médica, IPOC-FG, EPE, 3000-075 Coimbra, Portugal

joana@fe.uc.pt, hrocha@mat.uc.pt, brigida@ua.pt,

mclopes@ipocoimbra.min-saude.pt

Abstract. Radiation incidences (angles) that are used in Intensity Modulated Radiation Therapy (IMRT) treatments have a significant influence in the treatment clinical outcome. In clinical practice, the angles are usually chosen after a lengthy trial and error procedure that is significantly dependent on the planner's experience and time availability. The use of optimization models and algorithms can be an important contribution to the treatment planning, improving the quality of the solution reached and decreasing the time spent on the process. This paper describes a Dynamically Dimensioned Search (DDS) approach for IMRT beam angle optimization. Several different sets of parameters and search options were analyzed. Computational tests show that the final outcome is strongly influenced by these choices. This motivated the use of a cross-validation based procedure for choosing the algorithm's configuration, considering a set of ten retrospective treated cases of head-and-neck tumors at the Portuguese Institute of Oncology of Coimbra.

Keywords: Dynamically Dimensioned Search, IMRT, Beam Angle Optimization, Derivative-Free Methods

1 Introduction

Radiation therapy is one of the treatments used for cancer patients. Its aim is to destroy cancer cells through radiation, but at the same time spare healthy tissue that can also be damaged by radiation. The patient is usually immobilized on a couch that can rotate. The radiation is delivered through the use of a linear accelerator mounted on a gantry that can rotate along a central axis parallel to the couch. The rotation of the couch combined with the rotation of the gantry allows radiation from almost any angle around the tumor. Intensity Modulated Radiation Therapy is one type of radiation therapy where it is possible to modulate the radiation intensities that are delivered to the patient from each radiation incidence. This modulation is achieved through the use of a multileaf collimator. The collimator has left and right leaves that can block radia-

tion. By moving these leaves it is possible to create different intensity profiles (Fig. 1 and Fig. 2). Conceptually, this is equivalent to consider that, instead of having one single radiation beam from each radiation incidence used in the treatment, we can have a discretization of this beam into beamlets, each one with a given radiation intensity.




Fig. 1. Illustration of a multileaf collimator (with nine pairs of leaves)




Fig. 2. Illustration of a beamlet intensity map (9×9)

The possibility of modulating the radiation intensities increases the precision of the treatment and can be very important in diminishing treatment's side effects as it is possible to better spare cells that we do not want to irradiate. Nevertheless, it requires a complex planning procedure where many different and interconnected decisions have to be made by the planner, beginning by deciding how many and which angles to use in the treatment (the best angles can often be non-intuitive).

Whenever a patient is referred to an IMRT treatment, the medical doctor will delineate in the patient's computed tomographies (CT) the structures of interest: areas that should be treated plus a safety margin (Planning Target Volumes – PTV) and also the organs that should be spared (Organs at Risk - OAR). The medical doctor will also establish the medical prescription (Table 1), defining the desired dose for PTVs and mean or maximum doses to OARs (that sometimes are not possible to achieve). It is then up to the medical physicists to plan the treatment, by interacting with a Treatment Planning System (TPS) that simulates the behavior of the linear accelerator and calculates the radiation dose that will be deposited in the patient. By a trial and error procedure, several different treatment parameters are tried until a treatment plan that is considered admissible by the planner and that the planner thinks it is hard to improve is reached.

Table 1. Prescribed doses for all the structures considered for IMRT 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	–

Most of the times, the quality of the treatment plan is dependent on the planner's experience and time availability.

Another approach is to consider inverse planning, where the trial and error procedure is totally or partially replaced by the use of mathematical models and optimization algorithms. The mathematical models are characterized by being large-scale, non-linear and multi-modal, with objective functions that are computationally expensive to calculate. Global optimization algorithms that are derivative-based can easily be trapped in one of the many existing local minima.

In this paper we are concerned with the definition of the best radiation angles to use, considering the number of angles fixed *a priori* (Beam Angle Optimization – BAO). The BAO problem has been tackled using several different methodologies: scoring methods ([1]); methods based on the concept of beam's eye view ([2, 3, 4, 5]); response surface approaches ([6]); derivative-free approaches ([7]); mixed integer programming approaches ([8]); simulated annealing ([2, 9, 10]); particle swarm optimization ([11]); genetic algorithms ([12, 13, 14]) among others (see, for instance, [15, 16, 17, 18]).

In this paper, we consider an approach based on Dynamically Dimensioned Search (DDS). Several computational tests were done to understand the influence of the algorithm's configuration in the final outcome. The choice of the best configuration to use is not trivial, and we propose the use of a cross-validation procedure. This work follows a first preliminary experiment, where only one set of parameters was tested and that showed encouraging results [19].

The next section describes the mathematical model used. Section 3 describes the DDS algorithm. Computational results are shown and discussed in section 4. Section 5 states the main conclusions and presents future research ideas.

2 BAO Mathematical Model

The treatment planning process consists in determining, for a given patient, the angles that will be used in the treatment (BAO), the radiation intensities (fluences) for each of the angles (fluence map optimization - FMO), and the way that the multileaf collimator leaves should move to produce the desired fluence patterns (segmentation).

In this paper we are concerned with the BAO problem and we consider that the number of angles to use, k , will be fixed *a priori* by the planner. This means that we aim at finding out which is the best set of k angles out of every possible combination. For each beam angle set, we have to find a way of calculating the quality of this set. This can only be done after performing the fluence optimization, where the best radiation intensities for each of the angles considered will be calculated. To solve this FMO, the patient is discretized into voxels (small volume elements) and the radiation dose that is deposited in each of the patient's voxels is computed using the superposition principle, i.e., considering the contribution of each beamlet. Typically, a dose matrix D is constructed from the collection of all beamlet intensities, by indexing the rows of D to each voxel and the columns to each beamlet, i.e., the number of rows of matrix D equals the number of voxels (V) and the number of columns equals the

number of beamlets (N) from all beam directions considered. Therefore we can say that the total dose received by the voxel i is given by $\sum_{j=1}^N D_{ij} w_j$, with w_j the weight of beamlet j . Usually, the total number of voxels considered reaches the tens of thousands. If we define Θ as the set of all possible angles, then a basic formulation for the BAO problem can be defined as follows:

$$\min f(\theta_1, \theta_2, \dots, \theta_k) \quad (1)$$

$$\text{subject to } \theta_1, \dots, \theta_k \in \Theta \quad (2)$$

Many mathematical optimization models and algorithms have been proposed for the FMO problem and it is out of the scope of this paper to discuss the pros and cons of those models. In this paper, a convex penalty function voxel-based nonlinear model is used [20], such that, 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 quadratic programming problem with only linear nonnegativity constraints on the fluence values [21]. Let T_i be desired dose for voxel i , λ_i and $\bar{\lambda}_i$ the penalty weights of underdose and overdose of voxel i , respectively, and $(\bullet)_+ = \max\{0, \bullet\}$. Then the model can be defined as follows:

$$\text{Min}_w \sum_{i=1}^V \left[\lambda_i \left(T_i - \sum_{j=1}^N D_{ij} w_j \right)_+^2 + \bar{\lambda}_i \left(\sum_{j=1}^N D_{ij} w_j - T_i \right)_+^2 \right] \quad (3)$$

$$\text{s.t. } w_j \geq 0, j = 1, \dots, N \quad (4)$$

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 [21]. This nonlinear formulation implies that a very small amount of underdose or overdose may be accepted in clinical decision making, but larger deviations from the desired/allowed doses are decreasingly tolerated. Objective function (1) is calculated by (3), considering only beamlets that belong to $(\theta_1, \theta_2, \dots, \theta_k)$. This objective function is computationally expensive to calculate, taking up to a few minutes, depending on the patient itself and the number of angles considered.

3 DDS Algorithm

Considering the BAO problem, the DDS algorithm has as main advantages the fact that it is derivative-free, being able to escape from local minima, and the fact that it is possible to define *a priori* the number of objective function evaluations that will be

performed. This is especially important when dealing with a computationally expensive objective function.

The DDS algorithm begins with any admissible solution of the problem that becomes the current solution. In each iteration the algorithm finds a new solution by randomly perturbing the current one. Whenever a better solution is found, it becomes the current solution that, in turn, will be perturbed. The DDS algorithm can be interpreted as a random search process, considering searchable neighborhoods that will decrease in size as the algorithm iterates. This will promote a more global search at the beginning of the search and a more local search in the final iterations. In each iteration, each variable (angle) will be perturbed with a given probability. This probability decreases with the increase in the number of iterations, so that less and less angles are changed as the algorithm progresses. The magnitudes of the perturbations are randomly sampled from a normal distribution with mean 0. It is not necessary to consider an upper or lower bound for each variable, since an angle of -10° , for instance, is equal to 350° or an angle of 370° is equal to 10° . In our implementation of the algorithm we followed [22], considering some adaptations described in [23]. The algorithm's parameters are as follows:

- r_{init} represents the initial standard deviation considered;
- r_{max} and r_{min} represent the maximum and minimum admissible standard deviations considered;
- N represents the maximum number of iterations (an upper limit to the number of objective function evaluations, since in each iteration at most one solution is evaluated);
- $l_{success}$ and $l_{failure}$ determine a change in the current standard deviation due to successive successful or unsuccessful iterations (a success meaning that the objective function value has improved).

The algorithm has as input an admissible solution to the problem (that can be randomly generated) and returns as output an improved admissible solution (it is not possible to guarantee that it is optimal). The algorithm behavior can be described as follows:

1. Set counter $i \leftarrow 1$; Define the initial admissible solution $x_{current}$ and evaluate this solution ($f_{current}$). $f_{best} \leftarrow f_{current}$; $x_{best} \leftarrow x_{current}$; $success \leftarrow 0$; $failure \leftarrow 0$; $r \leftarrow r_{min}$.
2. Calculate the probability of any given variable (angle) be perturbed as $p(i) = 1 - \ln(i)/\ln(N)$. For each decision variable $x_{best}(j)$, $j=1, \dots, k$, add the variable to the set J with probability $p(i)$.
3. For every variable $x_{best}(j)$, $j \in J$, perturb randomly this variable considering a normal distribution $N(0, r)$. This perturbed solution will constitute the new $x_{current}$.
4. Evaluate $x_{current}$. If $f_{current} < f_{best}$, then $f_{best} \leftarrow f_{current}$; $x_{best} \leftarrow x_{current}$; $success \leftarrow success + 1$ and $failure \leftarrow 0$. Else $success \leftarrow 0$ and $failure \leftarrow failure + 1$.

5. If $failure \geq l_failure$ then $r \leftarrow \min(r/2, r_min)$.
6. If $success \geq l_success$ then $r \leftarrow \max(2r, r_max)$.
7. $i \leftarrow i + 1$. If $i \geq N$ then stop, else go to 2.

Steps 2 and 3 of the algorithm are responsible for calculating a new current solution in a random manner (by randomly deciding which angles to perturb and the magnitude of the perturbation). Given the specificities of the BAO problem, we also guarantee that the current solution does not have two adjacent angles that are too near each other. From a clinical point of view, angles that are less than 4° apart are considered the same. The evaluation of the current solution in step 4 is done by resorting to the optimization of the FMO problem, considering the angles defined by the current solution. Step 5 introduces a dynamic in the DDS algorithm considering that after some failed trials it is time to look for solutions in a narrower neighborhood, and when there are successful trials the searchable neighborhood can be wider.

4 Computational Experiments

The DDS algorithm was tested considering ten clinical examples of retrospective treated cases of head-and-neck tumors at the Portuguese Institute of Oncology of Coimbra (IPOC). A typical head-and-neck treatment plan consists of radiation delivered from 5 to 9 equally spaced coplanar orientations around the patient. The optimization of the angles has an increased importance when fewer angles are used. Being able to deliver a high quality treatment with fewer angles is beneficial for both the patient and the health institution. From the institution point of view, fewer angles mean faster treatment times, so that more patients can be treated. From the patient point of view, the faster the treatment the better because it is more likely that the patient does not change his position significantly during the treatment, which contributes to more accurate treatment results. For these reasons, treatments with 5 coplanar beams were considered.

In order to facilitate convenient access, visualization and analysis of patient treatment planning data, as well as dosimetric data input for treatment plan optimization research, the computational tools developed within MATLAB and CERR – computational environment for radiotherapy research ([24]) are used widely for IMRT treatment planning research. The ORART – operations research applications in radiation therapy ([25]) collaborative working group developed a series of software routines that allow access to influence matrices, which provide the necessary dosimetry data to perform optimization in IMRT. CERR was elected as the main software platform to embody our optimization research. 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 and MATLAB 7.4.0 (R2007a). The dose was computed using CERR’s pencil beam algorithm (QIB). For each of the ten head-and-neck cases, the sample rate used for Body was 32 and for the remaining structures was 4, resulting in 20,874 to 24,158 voxels and 948 to 1,283 beamlets for the 5-beam equispaced coplanar treatment plans. An automatized procedure for dose computation for each given beam angle set was

developed, instead of the traditional dose computation available from IMRTP module accessible from CERR's menubar. This automatization of the dose computation was essential for integration in our DDS algorithm. 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. For this set of patients, 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.

4.1 Clinical Examples

The selected clinical examples were signalized at IPOC as complex cases where proper target coverage and organ sparing, in particular parotid sparing, proved to be difficult to obtain. The patients' CT sets and delineated structures were exported via Dicom RT to CERR. Since the head-and-neck region is a complex area where, e.g., the parotid glands are usually in close proximity to or even overlapping with the target volume, careful selection of the radiation incidence directions can be determinant to obtain a satisfying treatment plan.

The spinal cord and the brainstem are some of the most critical organs at risk (OARs) in the head-and-neck tumor cases. These are serial organs, i.e., organs such that if only one subunit is damaged, the whole organ functionality is compromised. Therefore, if the tolerance dose is exceeded, it may result in functional damage to the whole organ. Thus, it is extremely important not to exceed the tolerance dose prescribed for these types of organs. Other than the spinal cord and the brainstem, the parotid glands are also important OARs. The parotid gland is the largest of the three salivary glands. A common complication due to parotid glands irradiation is xerostomia (the medical term for dry mouth due to lack of saliva). This decreases the quality of life of patients undergoing radiation therapy of head-and-neck, causing difficulties to swallow. The parotids are parallel organs, i.e., if a small volume of the organ is damaged, the rest of the organ functionality may not be affected. Their tolerance dose depends strongly on the fraction of the volume irradiated. Hence, if only a small fraction of the organ is irradiated the tolerance dose is much higher than if a larger fraction is irradiated. Thus, for these parallel structures, the organ mean dose is generally used instead of the maximum dose as an objective for inverse planning. In general, the head-and-neck region is a complex area to treat with radiotherapy due to the large number of sensitive organs in this region (e.g., eyes, mandible, larynx, oral cavity, etc.). 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. The parotid glands are in close proximity to or even overlapping with the PTV which helps explaining the difficulty of parotid sparing. Adequate beam directions can help on the overall optimization process and in particular in parotid sparing.

4.2 Results

For each BAO problem, the DDS algorithm was executed considering different configurations for the algorithm. The parameters that are expected to have a greater impact in the algorithm's outcome are the initial standard deviation (r_{init}), $l_{failure}$ and $l_{success}$. The last two are responsible for the evolution of the r parameter. However, after some preliminary tests, it was possible to conclude that there are seldom two consecutive successful iterations, so that if $l_{success}$ takes values greater than 1 this is equivalent to never changing r according to step 6 of the algorithm. Parameters r_{max} and r_{min} can and should be defined considering the specificities of the problem. In this case it was considered r_{max} equal to 90° and r_{min} equal to 3° . As we are randomly perturbing an angle using a normal distribution of mean 0 and standard deviation r , we know that there is 95% of probability of generating a perturbation value that belongs to the interval $[-2r, 2r]$. Notice that the greatest perturbation that is interesting to consider is 180° . Table 2 presents the values that were considered for parameters r_{init} and $l_{failure}$. Regarding $l_{success}$, it was considered to be fixed to 1 for the reasons exposed. The choice of the r_{init} values was motivated by the number of angles considered and the equidistant solution where all angles are 72° apart.

Table 2. Values of the Parameters

r_{init}	18	36	72
$l_{failure}$	5	20	

It is also interesting to consider the DDS algorithm when r stays constant throughout the algorithm's execution. This means that steps 5 and 6 are not considered.

The choice of increasing r in successive successful iterations and decreasing it after a sequence of failed iterations is an option that can be as justifiable as doing exactly the opposite. Notice that the algorithm's convergence is being guaranteed by the fact that the probability of perturbing each variable decreases iteration after iteration. So, we chose to also test this different version of the algorithm (Steps 5 and 6 will be replaced by Steps 5a and 6a).

5 a. If $failure \geq l_{failure}$ then $r = \max(2r, r_{max})$.

6 a. If $success \geq l_{success}$ then $r = \min(r/2, r_{min})$.

We have also tested a simpler rule, where r is randomly generated after $l_{failure}$ successive failed iterations (Steps 5 and 6 are replaced by Step 5b).

5 b. If $failure \geq l_{failure}$ then r is randomly generated using a uniform distribution in $[r_{min}, r_{max}]$.

For each of the ten patients, and for each version of the algorithm, five different runs were considered because of the random nature of the algorithm. A total of N equal to 200 iterations was considered. The initial solution considered was always the equidistant solution, as this is most of the times also the solution used in clinical prac-

tice. So, we are interested in measuring the improvement of the objective function value of the final solution (f_{DDS}) when compared with the equidistant initial one (f_{equi}). This improvement is calculated as $(f_{equi} - f_{DDS})/f_{equi}$. Before showing the global computational results, it is also worth to look at the influence of each parameter in the algorithm's behavior.

When the standard deviation r is kept constant, then we should expect a smooth behavior with smaller r_{init} values. Fig. 3 depicts the situation for a run of the algorithm considering patient 5. This patient was randomly selected, and similar behaviors are observed in the other patients.




Fig. 3. Algorithm's behavior for different r_{init} values

A similar behavior can be seen even when r_{init} is indeed only an initialization parameter. Smaller initial values are associated with smoother objective function values transitions. As the change in r is considered as dividing or multiplying its value by 2, the influence of r_{init} is present in all iterations. This can be seen in Fig. 4 and Fig. 5.




Fig. 4. Algorithm's behavior for different r_{init} values, $l_{failure}=5$




Fig. 5. Algorithm's behavior for different r_{init} values, $l_{failure}=20$

The impact of the $l_failure$ parameter is more visible with greater values of r_init . By inspection of Fig. 4 and Fig. 5, we can see that smaller values of $l_failure$ promote a faster convergence of the algorithm when r_init is equal to 36 or 72.

The option of using Step 5a and Step 6a allows a steepest descent in early iterations, but diminishes the successful search iterations as the algorithm progresses. This behavior is more pronounced for greater values of r_init .




Fig. 6. Algorithm's behavior considering Step 5a and Step 6a

The algorithm was run 5 times for each configuration considered. The BAO problem is characterized by having multiple local minima, so it is expected that in each run of the algorithm a different solution is found. This is illustrated in Fig. 7, where the equidistant solution is shown [black solid line] together with 5 other solutions that were calculated in each of the algorithm's runs.




Fig. 7. Different runs of the algorithm usually end up with different solutions

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

r_init	l_failure	Steps	Patients										Average improvement
			1	2	3	4	5	6	7	8	9	10	
18	-	-	2,75%	6,74%	7,49%	5,12%	7,28%	5,88%	14,65%	7,22%	6,82%	2,93%	6,69%
36	-	-	3,14%	7,05%	8,03%	4,38%	7,56%	5,59%	15,48%	8,33%	5,58%	2,42%	6,76%
72	-	-	3,24%	6,17%	8,67%	3,78%	7,88%	6,83%	14,16%	8,36%	5,43%	2,04%	6,66%
18	5	5, 6	3,00%	5,58%	7,22%	4,52%	7,77%	4,75%	14,94%	6,67%	5,54%	2,46%	6,25%
18	20	5, 6	3,24%	5,76%	7,71%	4,88%	8,15%	6,41%	14,40%	9,29%	6,63%	1,96%	6,84%
36	5	5, 6	3,22%	5,87%	8,74%	4,41%	7,85%	5,28%	14,43%	7,12%	5,92%	3,30%	6,61%
36	20	5, 6	3,25%	5,95%	8,82%	4,72%	7,64%	6,43%	14,57%	7,95%	5,95%	3,65%	6,89%
72	5	5, 6	2,71%	6,59%	7,32%	4,56%	6,94%	5,70%	15,31%	6,62%	5,10%	1,73%	6,26%
72	20	5, 6	2,93%	7,00%	7,03%	4,35%	7,99%	6,04%	13,29%	7,91%	6,17%	2,47%	6,52%
18	5	5a, 6a	3,05%	6,22%	7,17%	4,76%	7,32%	6,25%	16,81%	7,94%	5,39%	3,57%	6,85%
18	20	5a, 6a	3,19%	6,19%	8,92%	4,56%	7,46%	5,04%	16,44%	6,98%	5,20%	2,28%	6,63%
36	5	5a, 6a	3,06%	6,97%	7,39%	4,29%	7,56%	5,66%	15,71%	9,18%	5,25%	3,19%	6,82%
36	20	5a, 6a	3,27%	6,74%	6,90%	5,00%	7,20%	5,63%	14,77%	7,25%	6,51%	2,97%	6,62%
36	-	5b	3,22%	6,40%	7,23%	4,70%	8,37%	5,94%	14,52%	7,08%	5,74%	2,39%	6,56%
Average			3,09%	6,38%	7,76%	4,57%	7,64%	5,82%	14,96%	7,71%	5,80%	2,67%	
maximum			3,27%	7,05%	8,92%	5,12%	8,37%	6,83%	16,81%	9,29%	6,82%	3,65%	
minimum			2,71%	5,58%	6,90%	3,78%	6,94%	4,75%	13,29%	6,62%	5,10%	1,73%	

Table 3 shows the average improvement achieved in the objective function value. For each patient, the highest mean improvement obtained is highlighted. We see that the choice of the algorithm's configuration can have an important impact in the quality of the solution reached. There is no single configuration that appears as being the best one for a significant part of the patients: most algorithms are the best for one or two patients at the most. In clinical practice, due to time constraints, it is not possible to run the algorithm with different configurations and then choose the best solution reached. So, how should we decide which configuration to consider? One trivial choice would be to consider the one that would, on average, be the best one over all the patients tested. This approach can, however, be misleading.

The approach proposed in this paper is to consider cross-validation. This means that we select a set of patients, and with this set of patients all versions of the algorithm are ran. The best configuration, on average, for this set of patients, is then applied to the rest of the patients not belonging to this "cross-validation set". We have chosen *leave-one-out* cross-validation:

1. Select one patient j at a time. Consider a set constituted by all patients but j .
2. Run all versions of the algorithm, 5 times each. Calculate the mean improvement over all patients for each version of the algorithm.
3. Choose the version of the algorithm that presents the best objective value improvement. Apply this version of the algorithm to patient j , running the algorithm 5 times and recording the results.
4. Repeat the process for every available patient.

This *leave-one-out* cross-validation procedure can be implemented in a clinical setting, since the time constraints that exist are mainly concerned with guaranteeing that new patients are treated as soon as possible. So, it would be feasible to run several times the algorithm for each already treated patient, keeping a database with these results, and resorting to this database whenever it is necessary to choose a given version of the algorithm to apply to a new patient. The set of patients to include in this set could even consider some measures of similarity between patients.

Applying this procedure with our set of 10 patients, the results are as depicted in Table 4.

Table 4. Computational results when parameters are chosen by cross-validation

patient	r_{init}	$L_{failure}$	algorithm	$fequi$	mean $fDDS$	% improvement	Standard deviation
1	36	20	Steps 5,6	387,28	374,70	3,25%	1,07
2	36	20	Steps 5,7	72,93	68,59	5,95%	1,28
3	18	5	Step 5a, 6a	187,65	174,20	7,17%	3,61
4	36	20	Steps 5,6	156,37	148,99	4,72%	1,30
5	36	20	Steps 5,7	277,60	256,40	7,64%	2,25
6	36	5	Step 5a, 6a	165,58	156,21	5,66%	1,36
7	36	20	Steps 5,6	40,35	34,48	14,57%	0,72
8	36	20	Steps 5,7	166,08	152,87	7,95%	2,10
9	18	5	Step 5a, 6a	124,25	117,55	5,39%	1,65
10	18	20	Steps 5,7	186,44	182,77	1,96%	2,06

On average we are able to improve the objective function value 6,43%. For many other optimization problems, this would seem as a modest improvement. However, IMRT optimization problems have specificities that make the improvement in the objective function only one amongst several other criteria that can be used to assess the quality of the proposed optimization algorithm. More than the value of an objective function, the impact on the quality of the treatment for each patient is what really matters. The objective function is just a way of guiding the search for a better solution, but it cannot represent the whole set of complex features that have to be taken into account when assessing and considering admissible a given treatment plan.

A metric usually used for plan evaluation is the volume of PTV that receives 95% of the prescribed dose. Typically, 95% of the PTV volume is required as a minimum. These metrics are displayed for the ten cases in Fig. 8, considering the equidistant solution and the best and worst solutions out of the 5 solutions generated for each patient. The horizontal lines represent 95% of the prescribed dose. Satisfactory treatment plans should obtain results above these lines. By simple inspection we can verify the advantage of DDS treatment plans that have an improved tumor irradiation metric for most cases compared to equidistant treatment plans.

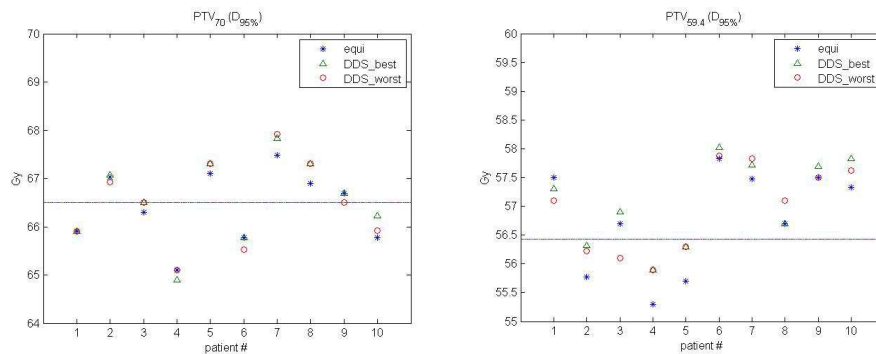


Fig. 8. Comparison of target irradiation metrics using DDS and equidistant treatment plans

In order to verify organ sparing, mean and/or maximum doses of OARs are usually displayed. These metrics are displayed for the ten cases in Fig. 9. The horizontal lines represent the tolerance mean or maximum dose for the corresponding structures. Satisfactory treatment plans should obtain results under these lines. For spinal cord, all treatment plans satisfy the maximum dose tolerance. For brainstem, treatment plans fulfill the maximum dose tolerance in almost all tested cases. Considering the mean dose limit for parotids, it was achieved less times. Looking at the right parotid, about half the patients receive an amount of radiation above what is desired. For the left parotid, the DDS optimized solutions guarantee a desirable level of radiation for 8 of the patients. Observing Fig. 9, it is perceivable that DDS treatment plans outperform equidistant treatment plans in terms of mean dose obtained.

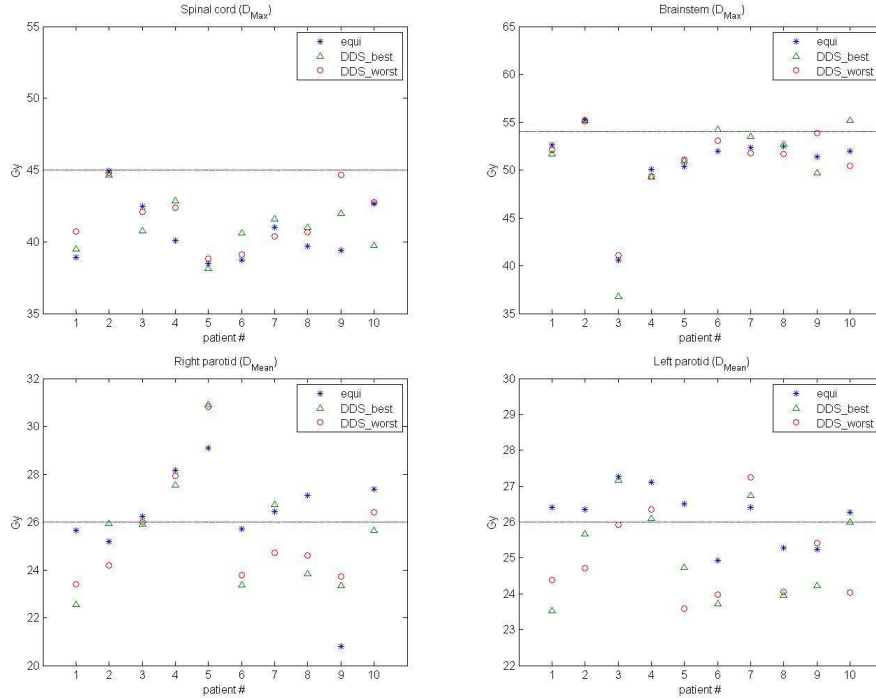


Fig. 9. Comparison of organ sparing metrics

Fig. 8 and Fig. 9 allow us also to illustrate the limitations of using a single objective function value to assess the quality of the solution. Looking at the results for patient 2, for instance, we can see that the solution that has the worst objective function value out of the 5 runs of the algorithm is, in fact, better when looking at the dose deposited in the patient. For some of these patients, namely those that are not getting a sufficient coverage of the PTV even with the optimized solutions, the next step would be to plan treatments with an increased number of radiation angles.

5 Conclusions

BAO problem is a very difficult global optimization problem, characterized by being a large, non-linear and multi-modal problem with a computationally expensive objective function. The DDS approach presented in this paper has as major advantages the fact that it is easily implemented, it is possible to determine the number of function evaluations that are performed and is a derivative-free search strategy that will not get easily trapped in a local minimum. Computational results show that the approach is capable of improving the equidistant solution. The calculation of optimized solutions are important not only contributing to the improvement of the treatment delivered to the patient considering the number of radiation incidences usually

determined *a priori*, but also allowing the planner to conclude that it will be necessary to increase the number of angles in order to reach an admissible treatment plan.

Further work will consider some changes in the proposed algorithm, namely embedding the DDS concept of neighborhood into a Simulated Annealing approach. It will also be necessary to consider the calculation of sets of solutions, instead of one single solution, that can illustrate the multiobjective inherent nature of this problem.

Acknowledgements. 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. The work of H. Rocha was supported by the European Social Fund and Portuguese Funds.

References

1. D'Souza, W., Meyer, R.R., Shi, L.: Selection of beam orientations in intensity-modulated radiation therapy using single-beam indices and integer programming. *Physics in medicine and biology*. 49, 3465-3481 (2004)
2. Lu, H.M., Kooy, H.M., Leber, Z.H., Ledoux, R.J.: Optimized beam planning for linear accelerator-based stereotactic radiosurgery. *International journal of radiation oncology, biology, physics*. 39, 1183-1189 (1997)
3. Goitein, M., Abrams, M., Rowell, D., Pollari, H., Wiles, J.: Multi-dimensional treatment planning: II. Beam's eye-view, back projection, and projection through CT sections. *International Journal of Radiation Oncology* Biology* Physics*. 9, 789-797 (1983)
4. Pugachev, A., Xing, L.: Pseudo beam's eye-view as applied to beam orientation selection in intensity-modulated radiation therapy. *International Journal of Radiation Oncology* Biology* Physics*. 51, 1361-1370 (2001)
5. Pugachev, A., Xing, L.: Computer-assisted selection of coplanar beam orientations in intensity-modulated radiation therapy. *Physics in Medicine and Biology*. 46, 2467-76 (2001)
6. Aleman, D., Romeijn, H., Dempsey, J. *A Response Surface-Based Approach to Beam Orientation Optimization in IMRT Treatment Planning*. in *IIE Annual Conf. Exposition*. 2006. Orlando, FL.
7. 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)
8. Lee, E.K., Fox, T., Crocker, I.: Simultaneous beam geometry and intensity map optimization in intensity-modulated radiation therapy. *International Journal of Radiation Oncology, Biology, Physics*. 64, 301-320 (2006)
9. Djajaputra, D., Wu, Q., Wu, Y., Mohan, R.: Algorithm and performance of a clinical IMRT beam-angle optimization system. *Physics in Medicine and Biology*. 48, 3191-3212 (2003)
10. Bortfeld, T., Schlegel, W.: Optimization of beam orientations in radiation therapy: some theoretical considerations. *Physics in Medicine and Biology*. 38, 291-304 (1993)
11. 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)

12. Dias, J., Rocha, H., Ferreira, B., Lopes, M.d.C.: A genetic algorithm with neural network fitness function evaluation for IMRT beam angle optimization. *Central European Journal of Operations Research*. accepted for publication, (2013)
13. Wu, X., Zhu, Y., Dai, J., Wang, Z.: Selection and determination of beam weights based on genetic algorithms for conformal radiotherapy treatment planning. *Physics in Medicine and Biology*. 45, 2547–58 (2000)
14. Li, Y., Yao, J., Yao, D.: Automatic beam angle selection in IMRT planning using genetic algorithm. *Physics in medicine and biology*. 49, 1915 (2004)
15. Ehr Gott, M., Johnston, R.: Optimisation of beam directions in intensity modulated radiation therapy planning. *OR Spectrum*. 25, 251-264 (2003)
16. 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)
17. Das, S.K., Marks, L.B.: Selection of coplanar or noncoplanar beams using three-dimensional optimization based on maximum beam separation and minimized nontarget irradiation. *International journal of radiation oncology, biology, physics*. 38, 643 (1997)
18. Craft, D.: Local beam angle optimization with linear programming and gradient search. *Physics in medicine and biology*. 52, N127-N135 (2007)
19. Dias, J., Rocha, H., Ferreira, B.C., Lopes, M.d.C. *IMRT Beam Angle Optimization Using Dynamically Dimensioned Search*. in *The International Conference on Health Informatics, IFMBE Proceedings*. 2013.
20. Aleman, D.M., Kumar, A., Ahuja, R.K., Romeijn, H.E., Dempsey, J.F.: Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning. *Journal of Global Optimization*. 42, 587-607 (2008)
21. Romeijn, H.E., Ahuja, R.K., Dempsey, J.F., Kumar, A., Li, J.G.: A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Physics in Medicine and Biology*. 48, 3521-3542 (2003)
22. Tolson, B., Shoemaker, C.: Dynamically dimensioned search algorithm for computationally efficient watershed model calibration. *Water Resources Research*. 43, W01413 (2007)
23. Regis, R.G., Shoemaker, C.A.: Combining radial basis function surrogates and dynamic coordinate search in high-dimensional expensive black-box optimization. *Engineering Optimization*. 45, 529-555 (2013)
24. Deasy, J.O., Blanco, A.I., Clark, V.H.: CERR: A computational environment for radiotherapy research. *Medical Physics*. 30, 979-985 (2003)
25. Deasy, J., Lee, E.K., Bortfeld, T., Langer, M., Zakarian, K., Alaly, J., Zhang, Y., Liu, H., Mohan, R., Ahuja, R.: A collaborative for radiation therapy treatment planning optimization research. *Annals of Operations Research*. 148, 55-63 (2006)