CEJOR
DOI 10.1007/s10100-013-0289-4

ORIGINAL PAPER

A genetic algorithm with neural network fitness function evaluation for IMRT beam angle optimization

Joana Dias · Humberto Rocha · Brígida Ferreira · Maria do Carmo Lopes

Abstract Intensity Modulated Radiotherapy Treatment (IMRT) is a technique used in the treatment of cancer, where the radiation beams are modulated by a multileaf collimator allowing the irradiation of the patient using non-uniform radiation fields from selected angles. Beam angle optimization consists in trying to find the best set of angles that should be used in IMRT planning. The choice of this set of angles is patient and pathology dependent and, in clinical practice, most of the times it is made using a trial and error procedure or simply using equidistantly distributed angles. In this paper we propose a genetic algorithm that aims at calculating good sets of angles in an automated way, given a predetermined number of angles. We consider the discretization of all possible angles in the interval [0°, 360°], and each individual is represented by a chromosome with 360 binary genes. As the calculation of a given individual's fitness is very expensive in terms of computational time, the genetic algorithm uses a neural network as a surrogate model to calculate the fitness of most of the individuals in the population. To explicitly consider the estimation error that can result from the use of this surrogate model, the fitness of each individual is represented by an interval of values and not by a single crisp value. The genetic algorithm is capable of finding improved solutions, when compared to the usual equidistant solution applied in clinical practice. The genetic algorithm will be described and computational results will be shown.

J. Dias
Faculdade de Economia, Universidade de Coimbra, Coimbra, Portugal

J. Dias (✉) · H. Rocha
Inesc-Coimbra, Coimbra, Portugal
e-mail: joana@fe.uc.pt

B. Ferreira · M. do Carmo Lopes
Serviço de Física Médica, IPOC-FG, EPE, Coimbra and Departamento de Física, Universidade de Aveiro, I3N, Aveiro, Portugal

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

Keywords Genetic algorithms · Radiotherapy · IMRT · Surrogate model · Neural networks

1 Introduction

The goal of radiation therapy is to deliver a dose of radiation to the cancerous region to sterilize the tumor minimizing the damages to the surrounding healthy organs and tissues. Radiation therapy is delivered with the patient immobilized on a couch that can rotate around a vertical axis. For most types of cancer, radiation therapy is administered 5 days each week for 5–8 weeks (Lim 2008). Using small radiation doses daily instead of few larger doses helps protect healthy tissues in the tumor region. Typically, high energy photon beam radiation is generated by a linear accelerator mounted on a gantry that can rotate along a central axis (Fig. 1) 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. The aim is to be able to plan a treatment that is in accordance with the medical prescription in terms of radiation dose distribution. Usually, the medical prescription will define prescribed doses to the target volumes (the regions that have to be irradiated), and mean or maximum tolerance doses to the organs at risk (the regions that should be spared). Despite the fact that almost all angles can be used in radiation delivery, the use of coplanar angles (without couch rotation) is the most usual option. This is a way to simplify an already complex problem, and the angles considered lie in the plane of the rotation of the gantry around the patient. Regardless the evidence presented in the literature that appropriate radiation beam incidence directions can lead to a plan's quality improvement (Das and Marks 1997; Liu et al. 2006), in clinical practice, most of the time, the number of beam angles is assumed to be defined *a priori* by the treatment planner and the beam directions are still manually selected by the treatment planner who relies mostly on his experience.

An important type of radiation therapy is intensity modulated radiation therapy (IMRT), where the radiation beams are modulated by a multileaf collimator. Multileaf collimators enable the transformation of the beam into a grid of smaller beamlets of independent intensities, as illustrated in Figs. 2 and 3. Here, we will consider IMRT optimization problems using coplanar angles and we will assume that the number of beam angles is defined *a priori* by the treatment planner.

The decision-making process regarding IMRT treatment planning can be conceptually understood as having three main steps.

1. Beam Angle Optimization (BAO): deciding what is the best number of beam angles and their directions. This means deciding how many times does the gantry stop, and where should it stop.
2. Fluence Map Optimization (FMO): deciding which are the best beamlet intensities for each gantry position. This means calculating the optimal radiation intensity profile that will be delivered to the patient every time the gantry stops, so that the medical prescription is satisfied.
3. Leaf Sequencing Problem: determining how the leaves of the multileaf collimator should move, so that the optimal beamlet intensities calculated in the previous step are, in fact, delivered.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

Fig. 1 Linear accelerator

Fig. 2 Illustration of an MLC (with nine pairs of leaves)

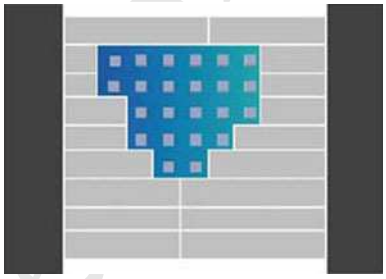

Fig. 3 Illustration of a beamlet intensity map (9 × 9)

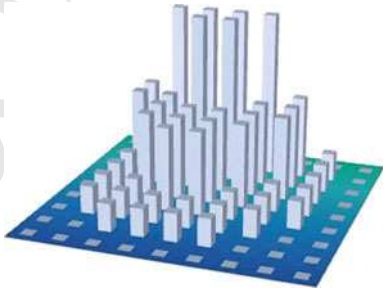

There are approaches that try to consider a beamlet-based approach, without the need to look at this problem using this three step framework, but those approaches lead to the development of large-scale optimization problems.

Most of the efforts in the IMRT optimization community have been devoted at optimizing beamlet intensities (Craft 2007), and comparatively fewer research effort has been directed to the optimization of beam angles (Ehrgott et al. 2008).

The BAO problem has been tackled using several different methodologies: scoring methods, where scores are assigned to beam angles based on geometric and dosimetric information (D'Souza et al. 2004); methods based on the concept of beam's eye view (Goitein et al. 1983; Lu et al. 1997; Pugachev and Xing 2001a,b), where the area of the tumor and the area of the surrounding organs as seen by the beam are considered in the selection of the better candidates; response surface approach (Aleman et al.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

2006); derivative-free approaches (Rocha et al. 2012); mixed integer programming approaches (Lee et al. 2006); among others (see, for instance, Das and Marks 1997; Ehrgott and Johnston 2003; Craft 2007; Lim and Cao 2012).

Many authors have also applied metaheuristics to this problem like simulated annealing (Bortfeld and Schlegel 1993; Lu et al. 1997; Djajaputra et al. 2003) or particle swarm optimization (Li et al. 2005). Evolutionary algorithms have also been used. Wu et al. (2000), consider conformal radiotherapy treatment planning, and use a genetic algorithm to determine beam intensities. Li et al. (2004), describe a genetic algorithm where each chromosome has as many genes as angles that are encoded using integers. Schreibmann et al. (2004), describe a hybrid multiobjective evolutionary algorithm that produces a set of efficient solutions for the multiobjective problem considered. Li and Yao (2006), use a genetic algorithm hybridized with an ant colony approach applied to BAO. Li et al. (2006), describe a genetic algorithm with a small population (the number of individuals is double the number of beams considered), that makes use of plan templates provided by experts. Bevilacqua et al. (2007), develop a genetic algorithm that can be applied to conformal, aperture-based and IMRT treatments, with the genetic representation changed accordingly. Lei and Li (2008), Li and Lei (2010), describe a genetic algorithm applied to the beam selection problem where the representation of each individual is based on the DNA structure. A small population of 24 individuals is used in chest and oropharyngeal tumor examples. Nazareth et al. (2009), describe the use of a genetic algorithm that is run on a distributed computing platform such that each generation takes about 30 min to be completed. Holdsworth et al. (2010), consider a two level multiobjective optimization evolutionary algorithm where the lower level performs a deterministic beamlet intensity optimization using a weighted quadratic objective function, and the top level uses a randomly generated population of individuals to represent the objective function weights that determine the relative weighting between organs to spare and targets. Fiege et al. (2011), describe the application of a parallel multiobjective genetic algorithm (*Ferret*) to BAO and FMO. They demonstrate the feasibility of the proposed approach on two phantoms (artificial models used to simulate the effect of radiation). Other genetic algorithm applications in radiotherapy treatment planning that do not deal explicitly with BAO can also be found (Lahanas et al. 2003; Haas and Reeves 2005).

The BAO problem is the first problem that should be solved in treatment planning, but in reality its optimal solution will be dependent on the optimal solutions of the two other sequential problems, especially on the optimal solution of the FMO. We need to know what are the optimal beamlet intensities associated with each set of angles to be able to compare different solutions (see, for instance, Das and Marks 1997; Haas et al. 1998; Schreibmann et al. 2004; Aleman et al. 2006; Craft 2007). When the BAO problem is not based on the optimal FMO solutions, the resulting beam angle set has no guarantee of optimality and has questionable reliability since it has been extensively reported that optimal beam angles for IMRT are often non-intuitive (Stein et al. 1997). Obtaining the optimal solution for a beam angle set is time costly and even if only one beam angle is changed in that set, a complete dose computation is required in order to compute and obtain the corresponding optimal FMO solution. This can cause a serious problem if we intend to use algorithms that require several objective function evaluations, as is the case with evolutionary algorithms, and we

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

have limited computational resources as is the case in many health institutions. One way of overcoming this problem is to use a surrogate model to calculate an estimation of the true objective function value. This idea has been used by several authors in different environments (see, for instance, El-Beltagy and Keane 1999; Emmerich et al. 2002; Jin and Sendhoff 2004; Jin and Branke 2005), but has never been applied to BAO.

The main contribution of this paper is to introduce the use of surrogate models within an evolutionary algorithm framework applied to BAO. We present a genetic algorithm that considers a discretization of the interval of all possible angles and that uses a neural network as surrogate model to calculate the fitness function of most individuals in the population. The algorithm uses the standard genetic operators (selection, crossover, mutation) as well as migration and a special local search operator that is nothing more than a genetic algorithm considering a very small elite population.

The paper is organized as follows: in the next section we describe the BAO problem formulation and the associated FMO problem formulation. Section 3 describes the surrogate model used. Section 4 describes the genetic algorithm. Clinical examples of head-and-neck cases used in the computational tests and some computational and clinical results are presented in Sect. 5. Section 6 presents the conclusions and some guidelines for future work.

2 Beam angle optimization problem

In beam angle optimization we aim at finding the best set of beams to be used in a given treatment. This means calculating the optimal number of beams, k, and figuring out what are the best k beam angles. This is a very important step in IMRT optimization since it directly influences both the quality of the treatment delivered and the overall treatment time. The treatment time increases with the increase in the number of beams. From the institution's point of view, fewer beams means 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.

In this paper we consider that k is determined *a priori*. For each set of k beams, we will need to determine a way of assessing the *goodness* of this set. This assessment can only be done after considering how the radiation dose will be deposited into the patient cells, so the FMO problem needs to be first solved so that we can consider the optimized beamlet intensities for each beam. To solve the FMO problem, we need a way to calculate accurately the radiation dose distribution deposited in the patient, measured in Gray (Gy). Each structure's volume is discretized in voxels (small volume elements) and the dose is computed for each voxel using the superposition principle, i.e., considering the contribution of each beamlet. Typically, a dose matrix D is such that each row of D corresponds to a voxel and each column to each possible beamlet. Thus, 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. The element in row i and column j of matrix D corresponds to the dose contribution to voxel i from beamlet j with unit intensity. Therefore we can say that the total dose

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

received by the voxel i is given by $\sum_{j=1}^{N} D_{ij} w_j$, with w_j representing the intensity (or fluence) of beamlet j. Usually, the total number of voxels considered reaches the tens of thousands, thus the row dimension of the dose matrix is of that magnitude. The size of D originates large-scale problems being one of the main reasons for the difficulty of solving the FMO problem. From a mathematical point of view, we are thus in the presence of two related problems. If we define Θ as the set of all possible angles, then a basic formulation for the BAO problem can be defined as follows:

$$\min f(\theta_1, \theta_2, \ldots, \theta_k) \tag{1}$$

$$\text{subject to } \theta_1, \ldots, \theta_k \in \Theta \tag{2}$$

If we consider a discretization of Θ then we are in the presence of a combinatorial optimization problem.

There is no consensual way of calculating optimal beamlet intensities. Many mathematical optimization models and algorithms have been proposed for the FMO problem, including linear models (Romeijn et al. 2003), mixed integer linear models (Lee et al. 2003), nonlinear models (Cheong et al. 2005), and multiobjective models (Craft et al. 2006). We have chosen to use a convex penalty function voxel-based nonlinear model (Aleman et al. 2008). According to this model, 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 nonegativity constraints on the fluence values (Romeijn et al. 2003):

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

$$\text{s.t. } w_j \geq 0, \quad j = 1, \ldots, N \tag{4}$$

where T_i is the desired dose for voxel i, λ_i and $\bar{\lambda}_i$ are the penalty weights of underdose and overdose of voxel i, respectively, and $(\cdot)_+ = \max\{0, \cdot\}$.

Although this formulation allows unique λ_i and $\bar{\lambda}_i$ weights associated with each voxel, similarly to the implementation in Aleman et al. (2008), 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 in clinical decision making, but larger deviations from the desired/allowed doses are decreasingly tolerated. It is beyond the scope of this study to discuss if this formulation of the FMO problem is preferable to others or not.

3 Surrogate model

For each set of k beam angles we will need to solve a FMO problem. The computational time spent solving one single instance of the FMO is dependent on the patient and

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

on k, but it is always computationally expensive. Table 2 shows the computational times in seconds of solving one instance of the FMO quadratic programming problem. Depending on the patient and number of beams considered, it can take from 40 s to more than 5 min to solve a single FMO optimization problem. This makes it difficult to use algorithms that require many objective function evaluations, like genetic algorithms.

To try and overcome this difficulty, it is possible to use a surrogate model, that will be able to estimate the true objective function value in a tiny fraction of time than it takes to calculate its true value.

In this work, we have decided to use a neural network (NN) to map sets of angles into objective function values. The use of NNs in radiotherapy problems is not new. Willoughby et al. (1996), describe the use of a NN to model the clinical judgment made by physicians when assessing treatment plans. Wells and Niederer (1998), develop an expert system based on NNs to plan standardized 3D conformal therapy treatments. Knowles et al. (1998), describe the use of NNs that will return the required treatment plan parameters after being trained (using evolutionary algorithms) with completed treatment plans. Rowbottom et al. (1999), consider prostate cancer patients treated with conformal therapy, and develop a NN that receives as inputs the patients' geometry information and has as outputs the angles' configuration. Gulliford et al. (2004), try to predict biological outcomes in prostate cancer patients after being submitted to radiotherapy treatments. Mathieu et al. (2005) and Vasseur et al. (2008), describe the use of NNs for accurate and fast dose calculation. Llacer et al. (2009), use NNs as part of an automatic non-coplanar beam orientation selection. The NNs work as classifiers, creating beam clusters based on the beam's coverage of the tumor to be treated plus some safety margins, and are also used for the geometric definition of the patient. Kalantzis et al. (2011), investigate the use of NNs in reconstructing dose maps for IMRT, achieving good results. For an introduction to NNs and a survey of NNs applied to radiotherapy problems see, for instance, Goodband and Haas (2008).

The main idea used in our work is as follows: we want to train a patient specific NN, that will receive sets of angles as inputs and that, for one particular patient, will return as output the value of the FMO objective function (3).

First of all, it will be necessary to generate a set of samples. Sets of k randomly generated angles are considered, and for each of these sets the true value of the objective function, f, is computed by calculating the optimal solution of the FMO problem. These samples are then used to train a neural network. The trained neural network is then ready to calculate f' expecting this value to be as close to f as possible (Fig. 4). This neural network should be capable of mapping a highly non-linear surface, with many local minima (Fig. 5). To facilitate the training of neural networks, and to improve their performance, both inputs and outputs are often normalized to lie in a fixed range (see, for instance, Shanker et al. 1996; Witten and Frank 2005; Han and Kamber 2006, pp. 70–72). We applied a normalization procedure as follows: considering the input angles, we decided to subtract 180° to each angle value and then divide by 180°, so that all angle values are within the interval $[-1,1]$. Regarding the output values, their statistical mean and standard deviation were calculated, the mean was subtracted from each value and the result was divided by the standard deviation. This results in a set of values with mean approximately equal to zero and standard deviation approximately equal to one.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

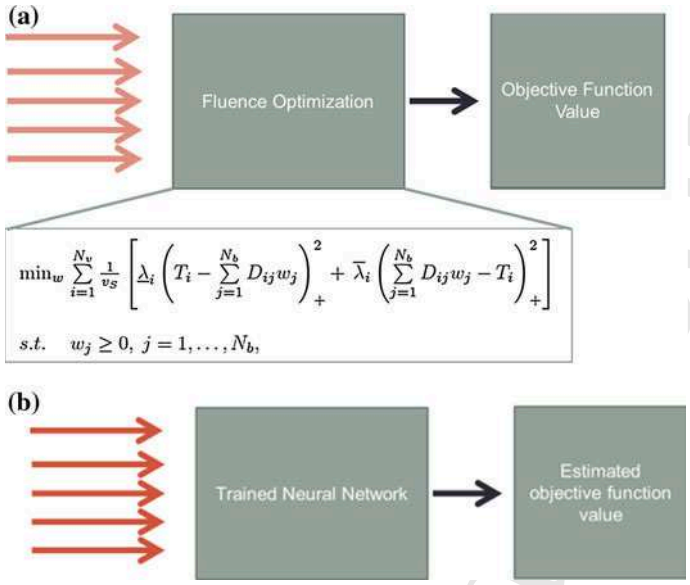

Fig. 4 Considering sets of 5 angles, the neural network will have five inputs (one input for each of the angles) and will have as output the estimated objective function value. **a** while training; **b** after training

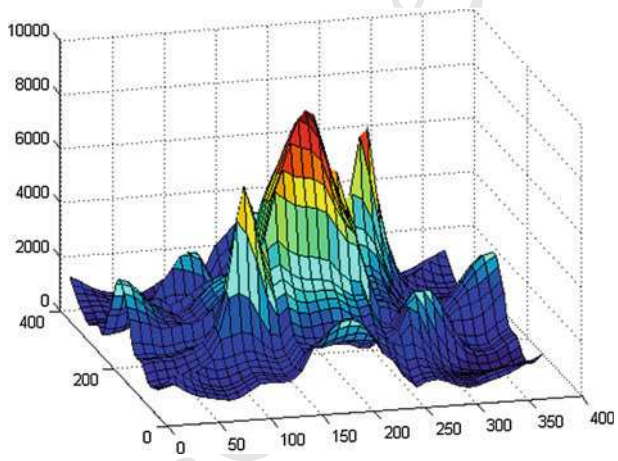

Fig. 5 Example of the objective function surface (defined by Eq. (3)) when we are considering sets of two angles only

It is not easy to decide on the best neural network architecture to use, and it is expected that this best architecture will be dependent on the particular situation at hand (the patient, the medical prescription, the number of angles, the number of samples available to train the NN).

For this reason, several neural network configurations are tested before choosing the best one. This is done as follows:

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

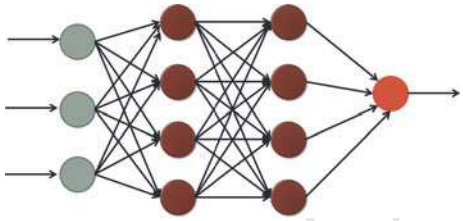

Fig. 6 An example of a neural network with three inputs (*three angles*), two hidden layers with four neurons each and one output (objective function value)

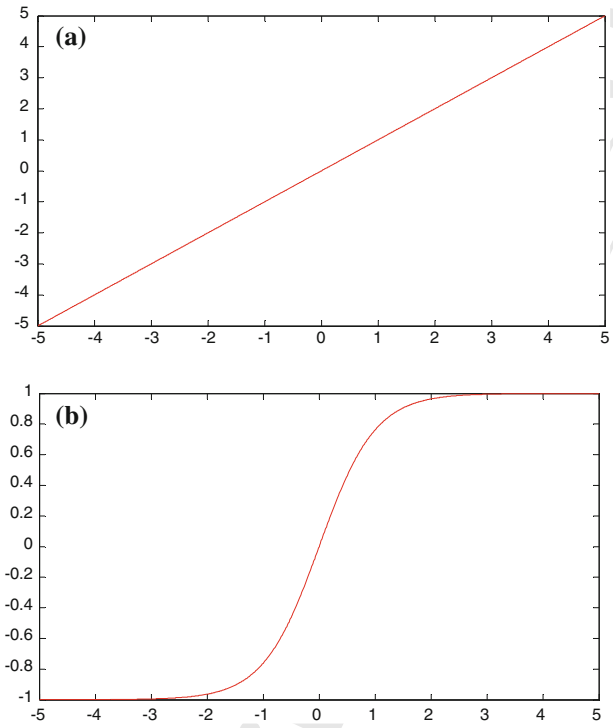

Fig. 7 **a** Linear transfer function; **b** Hyperbolic tangent sigmoid transfer function

1. Divide the set of available samples into three sets: training set (60 %), cross-validation set (20 %) and test set (20 %).
2. For each NN configuration, train the network using the samples in the training set, and calculate the estimation error using the cross-validation set. The estimation error is calculated by averaging the squared differences between each estimated value calculated by the NN and the corresponding target value from the cross-validation set.
3. Choose the NN configuration with the least estimation error.
4. Train 20 NNs with the same configuration chosen at step 3, randomly initialized.
5. Calculate the outputs for each of the 20 trained NNs using the test set, and consider the estimated values equal to the average value of all NN estimations. Calculate the expected estimation error using the test set.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

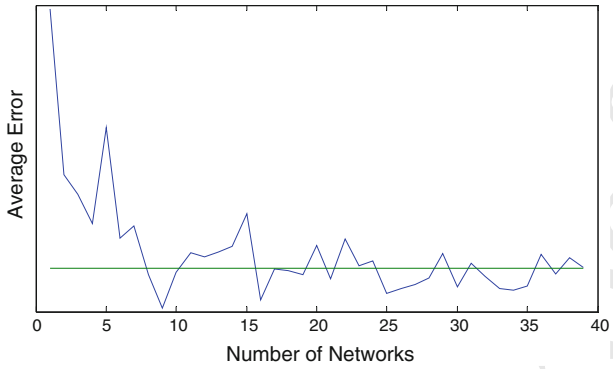

Fig. 8 Estimation error diminishes with the increase in the number of NNs

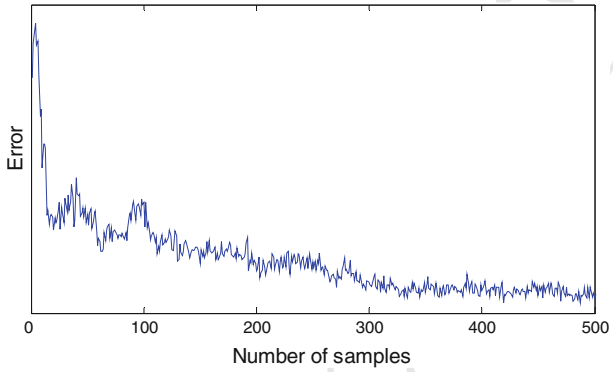

Fig. 9 Learning curve

The configurations tested in step 2 consider 1 to 5 hidden layers, and 1 to 40 neurons in each hidden layer (all hidden layers with the same number of neurons—Fig. 6). All hidden layers will have hyperbolic tangent sigmoid transfer functions, and the output layer will have a linear transfer function (Fig. 7).

The reasoning of using not a single NN but a set of 20 different NNs has to do with the fact that this can contribute to a decrease in the estimation error (Fig. 8).

As would also be expected, the error decreases with the increase in the number of available training samples (Fig. 9).

4 Genetic algorithm

The genetic algorithm developed considers the BAO problem as a combinatorial optimization problem, where the interval of all possible angles is discretized into 360 possible degrees. This makes it trivial to think of each individual (solution) as being represented by a chromosome constituted by 360 binary genes: if a gene is equal to one then the corresponding angle is used in the treatment, otherwise it is not. As the number of angles to be used is fixed *a priori* and is equal to k, this means that

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

each individual will have exactly k genes equal to one. The initial population will be randomly generated. As several sets of angles were already generated so that the NN could be trained, we take advantage of these individuals and use them as the initial population.

Two angles that are 4° or less apart from each other are considered similar from a clinical point of view. If a given individual has two angles that are less than 4° apart, then one of these angles is randomly chosen to be deleted (the corresponding gene is set to zero), and another randomly chosen gene is changed to 1. The procedure is repeated until there are no angles that are 4° or less apart. All individuals, in all generations, are submitted to such a procedure, so that we can guarantee that every individual will have angles that are at least 4° apart.

4.1 Fitness evaluation

For each individual in the population we can either calculate its fitness value by solving the FMO problem or by using the surrogate model. In the latter case, there will probably be an estimation error associated with the fitness value. This is why each individual will not have a single and crisp fitness value but its fitness will be represented by an interval. Each time the neural network is trained, we calculate the mean and the standard deviation of the estimation error using the test set. If we represent by μ the mean error and by σ the standard deviation, then if there were n samples in the test set and, for a given individual, the estimated fitness is given by f', then the individual's fitness interval will be equal to $[f' + \mu - 1.96\frac{\sigma}{\sqrt{n}}, f' + \mu + 1.96\frac{\sigma}{\sqrt{n}}]$. If the real fitness value is calculated using (3)–(4), then the upper and lower limit of this interval will be equal to the true fitness value.

It is not trivial to choose between using the original but computationally expensive objective function value or the surrogate model fitness calculation. It will always be necessary to calculate the true objective function value for some individuals, otherwise the function to be optimized will be the one represented by the surrogate model and not the true objective function. There are authors that consider the calculation of the true fitness for at least 50 % of all individuals, to guarantee that the evolutionary algorithm will try to optimize the true objective function (Hüsken et al. 2005). We can think about individual or generation control procedures (Jin et al. 2002). In the first approach, a certain number of individuals (*controlled individuals*) are evaluated using the true objective function in each generation. In the latter, in every g generations (*controlled generations*), all individuals are evaluated using the original fitness function. In this paper we have chosen the first approach, mainly due to computational limitations (the latter is better used when it is possible to use parallelization).

In each generation, two individuals are chosen and are passed directly to the next generation (elitist approach): the one that has the minimum fitness interval lower bound and the one that has the minimum fitness interval upper bound. Their true fitness function is calculated (if not known already). The choice of, in each generation, calculating the true objective value for two individuals, has to do with the fact that we need to assure the genetic algorithm is trying to reach an optimal solution for the "true" objective function, and is not optimizing the objective function defined by the surrogate model.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

It was tried to maintain a lookup table, that would keep a record of every individual with a known true objective fitness function. Then, whenever a new individual was generated, the first thing to do would be to check if there is a match in the lookup table. Nevertheless, a match occurred in less than 0.5 % of the times, so this option was abandoned.

4.2 Selection operator

The selection operator is responsible for the selection of individuals that will generate the offsprings that will constitute the new generation. We have chosen to use the tournament selection operator. This means that two individuals are randomly chosen from the current population, and their fitness values are compared. The best individual is chosen with a given probability. The procedure is repeated so that a second individual is chosen. These two individuals are then used to generate two new individuals.

As we are considering fitness intervals and not crisp fitness values, if the intervals do not overlap, then the comparison between individuals is trivial. If we are in the presence of two individuals such that their fitness intervals are, respectively, $\left[f_1' + \mu - 1.96\frac{\sigma}{\sqrt{n}}, f_1' + \mu + 1.96\frac{\sigma}{\sqrt{n}}\right]$ and $\left[f_2' + \mu - 1.96\frac{\sigma}{\sqrt{n}}, f_2' + \mu + 1.96\frac{\sigma}{\sqrt{n}}\right]$ such that $f_1' < f_2'$, then if $\left(f_1' + \mu + 1.96\frac{\sigma}{\sqrt{n}}\right) < \left(f_2' + \mu - 1.96\frac{\sigma}{\sqrt{n}}\right)$ we can conclude that individual 1 has a better fitness than individual 2. If this does not happen, meaning that $\left(f_1' + \mu + 1.96\frac{\sigma}{\sqrt{n}}\right) \geq \left(f_2' + \mu - 1.96\frac{\sigma}{\sqrt{n}}\right)$, then the fitness intervals overlap and we have to see if it is possible to conclude whether the fitness values are comparable or not (Knezevic 2008). If $\left(f_2' - f_1'\right) > 1.96\frac{\sigma}{\sqrt{n}}\sqrt{2}$ we consider that individual 1 has a better fitness than individual 2, otherwise we consider that we cannot reach a conclusion and we choose randomly one of them to participate in the crossover operator.

4.3 Crossover operator

The crossover operator used is as follows: each pair of parents will generate two twins. Each twin will have all angles belonging to each of the parents (Fig. 10). As we must have k and only k genes equal to one, these twins will correspond to non-admissible solutions (unless both parents are identical). So, a random procedure is applied, where genes that are equal to one are randomly selected and are changed to zero until k angles are reached.

Other crossover operators were tried (like one-point and two-point crossover), but better results were obtained with this procedure.

4.4 Mutation operator

Each offspring will, with a given probability, suffer a mutation. This means that one randomly chosen gene that is equal to one will be changed to zero, and one randomly chosen gene equal to zero will be changed to one.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

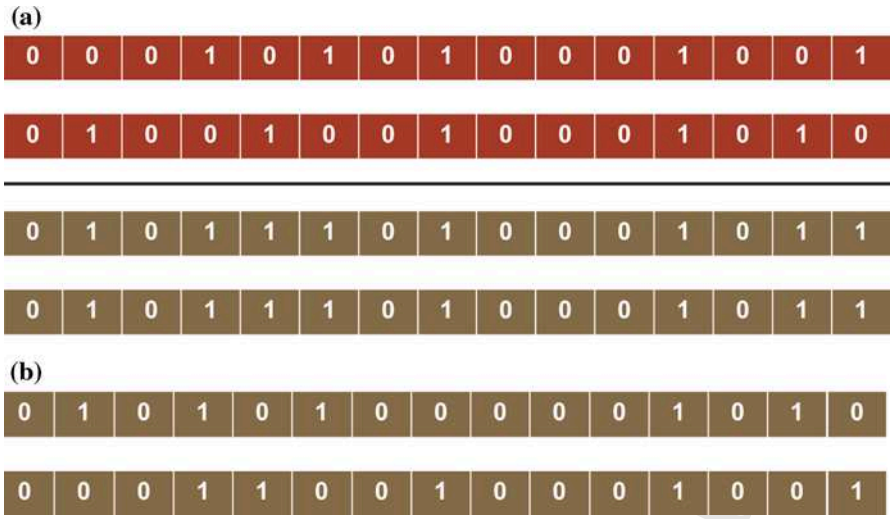

Fig. 10 Crossover operator **a** the two parents above will generate two twins. **b** these twins will be randomly changed so that only k genes keep their values equal to 1

4.5 Migration

Whenever some number of generations is evolved without an improvement in the objective function value of the BAO problem, a certain percentage of the population is substituted by randomly generated individuals. This procedure can be interpreted as a migration using a population with a high mutation rate, and contributes to the increase of the population diversity.

4.6 Local search

After running the described genetic algorithm, a local search procedure is executed. This local search procedure is, in fact, another genetic algorithm composed by exactly the same procedures of the genetic algorithm described, but with two main differences: the population is constituted by a very small number of individuals, and it is initialized by considering mutations of the best individual found so far; the fitness of all the individuals in this elite population is calculated by solving the FMO problem. Due to computational time limitations, the population will only evolve during a small number of generations.

4.7 Retraining the NN

In every generation of the genetic algorithm, the true objective function is calculated for the best individuals in the population. This means that these individuals can be considered as new samples. From time to time, the NN is retrained using this new and enlarged sample set. As the number of training samples increases, the standard deviation and the average error are expect to decrease, so that as the genetic algorithm evolves better estimations are produced by the surrogate model.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

4.8 The whole picture

The complete algorithm is now described:

1. Generate a set of samples, by randomly generating k angles and calculating the corresponding FMO objective function value (3).
2. Find the best NN architecture.
3. Train a set of NNs.
4. Execute the genetic algorithm
   a. Initialize the population
   b. While the termination condition is not met (maximum time or maximum number of iterations has not been reached)
      i. The true fitness value is calculated for two "best" individuals, that are immediately passed on to the next generation
      ii. Selection
      iii. Crossover
      iv. Mutation
      v. If the objective function does not improve during n consecutive generations then migration
      vi. If m new samples have been created, retrain the NNs
5. Execute the local search

5 Computational and clinical results

The described algorithm was tested considering ten clinical examples of already treated patient 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. We considered treatments with 5 coplanar beams because the importance of beam angle selection increases when a lower number of beam angles is considered. Furthermore, 5 angles is the usual starting point for the trial and error procedure conducted by planners. An increase in the number of angles is only considered if they are unable to reach a clinically acceptable solution.

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 (Deasy et al. 2003) are used widely for IMRT treatment planning research. The ORART—operations research applications in radiation therapy (Deasy et al. 2006) collaborative working group developed a series of software routines that 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 voxel size considered was 0.3 cm $\times$ 0.3 cm $\times$ 0.3 cm. Table 1

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

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

Structure	Spinal cord	Brainstem	Left parotid	Right parotid	PTV1	PTV2	Body
Mean dose	–	–	26 Gy	26 Gy	–	–	–
Maximum dose	45 Gy	54 Gy	–	–	–	–	80 Gy
Prescribed dose	–	–	–	–	70 Gy	59.4 Gy	–
# voxels case 1	1,382	1,715	1,576	1,390	4,001	31,119	1,790,592
# voxels case 2	3,567	2,072	1,536	1,807	1,485	43,649	1, 413,138
# voxels case 3	3,265	2,087	2,538	2,367	16,860	69,748	1,608,589
# voxels case 4	1,424	1,569	676	684	4,856	28,721	6,64,886
# voxels case 5	1,115	983	1,372	1,265	31,924	2,292	2,073,296
# voxels case 6	1, 101	1,518	1,176	1,140	30,047	6,613	1,560,070
# voxels case 7	985	851	668	631	19,835	3,973	1,110,882
# voxels case 8	1,160	1,223	1,405	1,323	29,786	4,450	1,710,982
# voxels case 9	829	1,135	782	1,096	11,348	1,153	1,016,083
# voxels case 10	533	2,907	1,056	662	25,461	11,066	1,553,317

presents the number of voxels for each patient and for each structure considered. An automated procedure for dose computation for each given beam angle set was developed, instead of the traditional dose computation available from CERR's menu bar. This automation of the dose computation was essential for integration in our BAO 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 to 350 s to be calculated, depending on the patient and on the set of beam angles considered (Table 2).

5.1 Clinical examples

Ten clinical examples of already treated patient cases of head-and-neck tumors at the Portuguese Institute of Oncology of Coimbra are used to test the genetic algorithm described. 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 (see Fig. 11). Since the head-and-neck region is a complex area where, e.g., the parotid glands (the two largest salivary glands) are usually in close proximity to or even overlapping with the target volume, careful selection of the radiation directions can be essential 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 impor-

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

Table 2 Computational times (in seconds) needed to solve one instance of the FMO quadratic programming problem (considering a sample of 100 instances for each patient)

Patient	5 angles			7 angles			9 angles		
	Min.	Average	Max.	Min.	Average	Max.	Min.	Average	Max.
1	43.20	94.30	102.20	126.51	130.89	137.33	98.48	185.71	226.22
2	56.44	66.17	75.96	81.01	97.61	113.85	136.79	106.98	218.36
3	50.85	117.02	126.70	157.94	169.38	174.12	133.54	235.30	255.91
4	45.68	110.45	122.63	71.00	89.98	101.33	107.55	124.70	142.12
5	89.17	96.64	113.85	122.06	139.38	153.36	170.47	191.95	236.05
6	71.07	82.73	96.16	99.38	121.32	206.72	131.46	157.21	250.99
7	59.20	77.34	91.97	97.35	123.39	152.48	140.24	186.24	227.21
8	78.26	95.33	105.46	124.21	138.98	158.59	157.84	191.27	323.66
9	73.39	91.98	194.84	102.75	128.98	164.13	148.03	185.64	350.94
10	83.60	93.62	103.00	119.23	131.40	147.64	162.98	185.62	228.50

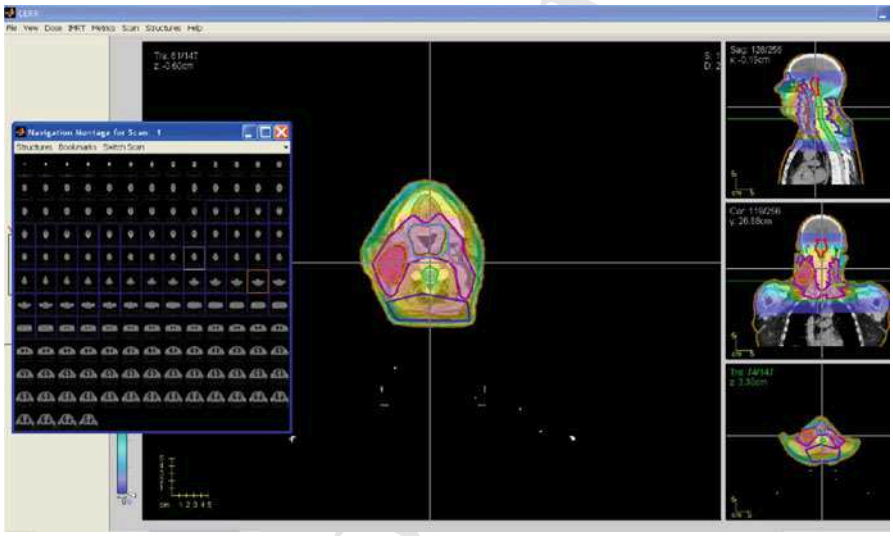

Fig. 11 CERR environment

tant OARs. The parotid glands are the largest of the three salivary glands. A common complication due to the irradiation of parotid glands 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 planning. In general, the head-and-

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

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.). In this study, the OARs used for treatment optimization were defined by the medical physicists as being the spinal cord, the brainstem and the parotid glands. The tumor to be treated plus some safety margins is called planning target volume (PTV). For the head-and-neck cases in study it was separated in two parts with different prescribed doses: PTV1 and PTV2. The prescription dose for the target volumes and tolerance doses for the organs at risk 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 sparing them. Adequate beam directions are an integral and important part of IMRT optimization, and can be determinant for achieving the sparing of parotid glands.

5.2 Computational results for NN

Neural networks were implemented by using the Matlab Neural Network Toolbox. The surrogate model was tested considering the ten different patients. To assess the behavior of the model in different settings, we considered five, seven and nine angles. It is not trivial to determine how should the performance of a surrogate model be measured, especially when this surrogate model is being used to guide an evolutionary algorithm (see, for instance, Hüsken et al. 2005). In this paper we will analyze the results obtained by considering the relative estimation error, especially looking at its average value and standard deviation. The results shown in Table 3 consider a training set of 100 samples. As can be observed, the error standard deviation decreases with the increase in the number of angles. This means that it will be possible to use fewer samples when dealing with more angles, without deteriorating in a significant way the quality of the estimation, as can be seen by looking at Table 4 that considers a training set of 50 samples.

If we try to fit a probability distribution to the estimation errors obtained, most of the times the normal and the logistic distributions are the most adequate considering the Anderson-Darling statistic.[1]

5.3 Computational results for GA

In the computational tests for GA we will consider IMRT treatments with five angles.

The genetic algorithm was implemented considering an initial population of 100 individuals (the individuals used to train the initial set of neural networks). The selection operator will choose the best individual with 80 % probability. We chose to use a high mutation rate of 50 %. Whenever 25 generations evolve without an improvement in the objective function value, 25 % of the population is replaced by randomly generated individuals. The local search procedure considers an initial population of only 10 individuals, created by mutating the best individual found so far, and this population is evolved during at most 10 generations.

[1] These tests were performed using the fit distribution option of software @Risk.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

Table 3 Computational results for NN (N-Number of neurons in each level; L-number of hidden layers; SD-relative error standard deviation; A-average relative error), 100 samples

Patient	5 angles				7 angles				9 angles			
	Best NN configuration		Relative error		Best NN configuration		Relative error		Best NN configuration		Relative error	
	N	L	SD (%)	A (%)	N	L	SD (%)	A (%)	N	L	SD (%)	A (%)
1	25	2	14	−2	8	5	9	−1	17	1	7	0
2	33	2	17	2	18	4	12	−1	33	2	8	0
3	20	3	9	0	37	2	6	−1	40	2	6	0
4	21	2	11	0	36	1	6	1	25	1	5	0
5	19	3	10	0	38	1	6	−1	37	2	4	0
6	29	2	15	0	22	2	9	2	30	4	6	0
7	40	2	25	0	16	4	20	9	36	2	13	−1
8	33	2	14	0	39	3	14	−1	35	3	10	0
9	23	2	15	0	28	3	12	0	35	2	8	−1
10	30	2	12	3	26	2	8	1	22	5	6	0

Table 4 Computational results for NN (N-Number of neurons in each level; L-number of hidden layers; SD-relative error standard deviation; A-average relative error), 50 samples

Patient	5 angles				7 angles				9 angles			
	Best NN configuration		Relative error		Best NN configuration		Relative error		Best NN configuration		Relative error	
	N	L	SD(%)	A (%)	N	L	SD (%)	A (%)	N	L	SD (%)	A (%)
1	27	2	15	−3	12	4	9	−1	3	4	7	1
2	40	2	20	3	12	3	13	0	10	2	8	1
3	22	5	9	−2	25	1	7	−2	28	2	6	−1
4	10	5	9	1	38	2	7	−1	9	2	5	0
5	16	4	12	0	36	2	6	−2	32	3	5	0
6	31	1	18	−1	5	4	10	0	17	5	6	−1
7	33	2	33	5	1	4	22	14	11	5	13	−3
8	33	1	19	6	33	3	15	−1	30	2	11	0
9	19	4	18	0	40	3	12	1	17	2	9	0
10	29	2	14	4	24	1	8	−2	3	4	7	−1

If we want to be able to apply these procedures in a clinical setting, then we have to consider some time constraints. It is important that planning the treatment of a given patient does not take more than one night. This means that we should consider as a time limit more or less 12 h. That is why we have chosen to terminate the genetic algorithm when 400 generations are reached or if 10 h have elapsed, whatever occurs sooner (notice that the time starts counting with the generation of the sampling sets used to train the NN). Then we allow the local search procedure to take at most 2 h.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

Table 5 Computational results for GA

Patient	Average *BAO-GA* solution	Standard deviation	*Equi* solution (%)
1	326.98	1.53	4.3
2	66.78	0.01	6.4
3	174.52	0.81	9.0
4	138.28	1.61	8.7
5	244.49	1.23	9.0
6	152.41	0.96	9.8
7	30.86	0.44	10.5
8	134.16	0.20	10.6
9	95.82	0.51	10.7
10	152.84	0.95	6.2

Table 6 Generating random solutions: computational results

Patient	Best random solution	Standard deviation	*Equi* solution (%)
1	336.57	5.81	1.5
2	67.70	3.25	5.1
3	182.66	3.28	4.8
4	141.78	2.73	6.4
5	261.83	1.66	2.5
6	162.95	1.37	3.5
7	34.70	1.26	0.0
8	151.10	1.34	0.0
9	103.49	1.00	3.5
10	162.21	4.22	0.5

The results of BAO optimization concerning the improvement of the objective function value for the ten cases of head-and-neck tumors using our BAO algorithm, denoted *BAO-GA* are presented in Table 5. Due to the random nature of the genetic algorithm, it is not sufficient to present results considering a single execution of the algorithm. We chose to execute the algorithm five times for each patient, and we present average and standard deviation results. The fourth column presents the average decrease in the objective function value when compared with the traditional 5-beam equispaced coplanar treatment plans, denoted *equi*.

Using the surrogate model implies the random generation of a set of solutions. It could be interesting to compare the results obtained by using the genetic algorithm with the results obtained using a simple random generation procedure. To be able to draw conclusions, it would not be advisable to compare the results of the GA with a single set of random solutions. As calculating the objective function value of each solution is very time consuming, we decided to randomly generate and evaluate 300 solutions. This is our base set. Then, using this set, we consider a random withdrawal of 100 solutions and calculate the best solution among the 100 solutions. This process is then repeated 100 times (so that in each time we get a possibly different sample of 100 solutions randomly taken from the base set). Table 6 presents the best solution found, the standard deviation calculated, and the improvement regarding the *equi* solution.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

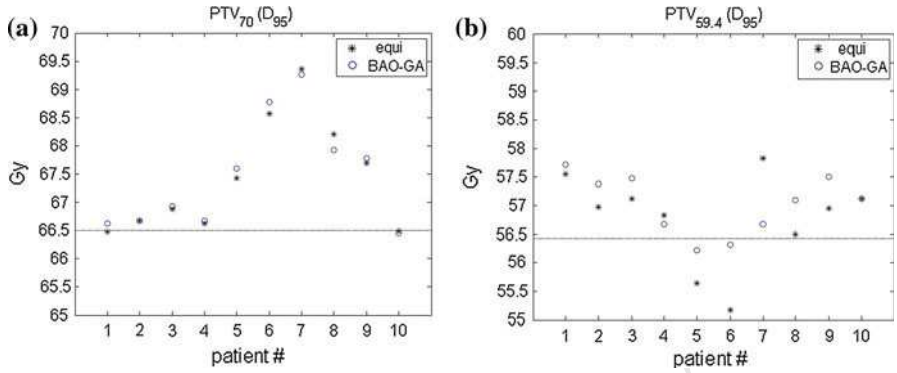

Fig. 12 Comparison of target irradiation metric obtained by *BAO-GA* and *equi* treatment plans

As can be seen, the genetic algorithm presents better results, and also a more reliable behavior, since the standard deviations are considerable when applying only a random procedure.

Since a small standard deviation was obtained for the results of the different runs of the *BAO-GA* algorithm, for the remainder of this section we will use the treatment plans corresponding to the best *BAO-GA* solution.

Despite the improvement in FMO value, the quality of the results can be perceived considering a variety of metrics. 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. This metric is displayed for the ten cases in Fig. 12. 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 *BAO-GA* treatment plans that have an improved tumor irradiation metric for most cases compared to *equi* treatment plans.

In order to verify organ sparing, mean and/or maximum doses of OARs are usually displayed. For each OAR, the corresponding metric is displayed for the ten cases in Fig. 13. 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 and brainstem, treatment plans fulfill the maximum dose tolerance in all tested cases. However, as expected, the mean dose limit for parotids was only achieved few times mostly by *BAO-GA* treatment plans. Moreover, observing Fig. 13, it is perceivable that *BAO-GA* treatment plans outperform *equi* treatment plans in terms of mean dose obtained. In fact, in average, *BAO-GA* treatment plans reduced the mean dose of the parotid glands by 0.96 Gy compared to the *equi* treatment plans.

Typically, results are judged by their cumulative dose-volume histogram (DVH). The DVH displays the fraction of a structure's volume that receives at least a given dose. DVH results for the sixth patient illustrate the numbers presented in Fig. 14. Since parotids are the most difficult organs to spare, as shown in Fig. 13, for clarity, the DVHs only include the targets and the parotids and were split in left and right parotid. The asterisks indicate 95 % of PTV volumes versus 95 % of the prescribed

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

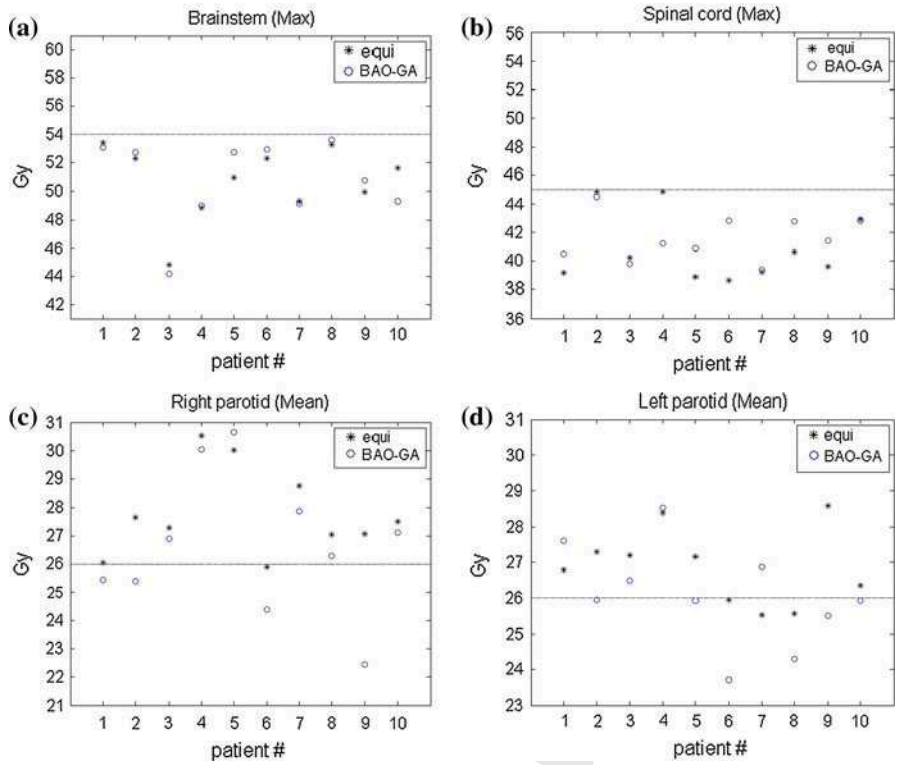

Fig. 13 Comparison of organ sparing metrics obtained by *BAO-GA* and *equi* treatment plans

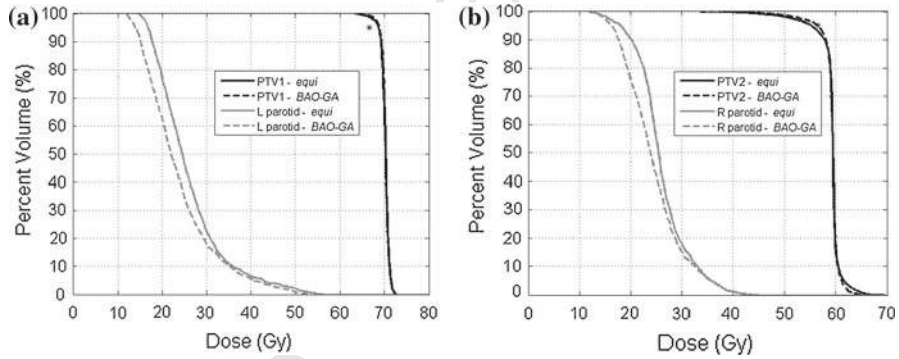

Fig. 14 Cumulative dose volume histogram comparing the results obtained by *BAO-GA* and *equi* treatment plans for the sixth patient

doses. The results displayed in Fig. 14 confirm the benefits of using the optimized beam directions, in particular using the directions obtained and used in the *BAO-GA* treatment plan.

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

6 Conclusions and future research

Beam angle optimization in radiotherapy treatment planning, especially in IMRT, can lead to significant improvements in the quality of treatments delivered to patients. It can lead to better preservation of the organs at risk, without jeopardizing the treatment efficacy, leading to an increase in the quality of life of patients. The need for beam angle optimization increases with the decrease in the number of angles to be used in a given treatment. Although better results can be expected with an increase in the number of directions used, using fewer angles is beneficial not only for the patient but also from the health institution's point of view: fewer angles means faster treatments, so more patients can be treated; faster treatments means better treatment precision because the probability of maintaining the patient immobilized in the desired position with no significant intrafraction setup errors (errors caused by organ motion or patient position change during treatment) is increased.

In this paper we introduce the use of patient dependent surrogate models embedded into an evolutionary algorithm optimization framework for BAO. The BAO problem is usually characterized by the existence of many local minima, and a highly nonlinear optimization surface. Genetic algorithms or evolutionary algorithms in general, are known to be able to tackle this kind of problems, due to their diversification capabilities and ability to escape from local minima. Nevertheless, due to the computational time needed to assess each given individual (solution), that can take from 1 min to more than 5 min, the use of genetic algorithms may not be compatible with clinical practice, especially with limited availability of computational resources. In clinical practice we should take no more than 12 h to generate an improved treatment. To try and overpass this difficulty, we propose the use of a surrogate model (a trained neural network) that will be used to calculate the fitness of most individuals in the population.

The computational results obtained show that the use of surrogate models combined with genetic algorithms can be an interesting path of research to follow.

Regarding future work, we will consider not only the improvement of the genetic algorithm, but also try to improve the surrogate model. In the latter case, instead of considering as inputs the angles, we can consider using scores associated with these angles (see, for instance, Pugachev and Xing 2002). The neural network will then receive information that is expected to be more related with the objective function value than only the angles. Another possibility is to consider a neural network that, instead of calculating an estimate of the objective function value, will be able to compare two individuals stating if one individual is better or worse than the other. As a matter of fact, in the genetic algorithm evolution, the selection procedure drives the evolution process, and more than knowing precisely the objective function values we need a way of comparing individuals in a fast and reliable way (it is more important to ensure the correct selection than to reproduce exactly the true fitness values—Hüsken et al. 2005). Another change would be to consider diminishing the number of times the real fitness function is calculated as the genetic algorithm evolves, as we will be dealing with a surrogate model that is improved whenever the neural networks are retrained.

In this paper we consider a discretization of the interval $[0°, 360°]$ in 360 values. We now feel that this may be too ambitious. New experiments will be made considering

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

at first a more coarse discretization, and only in later stages considering a larger set of discretized values.

Looking at the FMO problem, and thinking of an evolutionary algorithm, it is almost inevitable to think of multiobjective approaches to deal with this problem that is multiobjective by nature: we want to give the prescribed radiation to the target volumes, and as little radiation as possible to the organs at risk, which in most cases are contradictory objectives.

Acknowledgments This work was supported by FEDER funds through the COMPETE program, by QREN under Mais Centro (CENTRO-07-0224-FEDER-002003), and Portuguese funds through FCT—Fundação para a Ciência e a Tecnologia, under project PTDC/EIA-CCO/121450/2010 and by FCT under project grant PEst-C/EEI/UI0308/2011. The work of H. Rocha was supported by the European social fund and Portuguese funds.

References

Aleman D, Romeijn H, Dempsey J (2006) A response surface-based approach to beam orientation optimization in IMRT treatment planning. IIEAnnual conference exposition, Orlando, FL, pp 6–11

Aleman DM, Kumar A, Ahuja RK, Romeijn HE, Dempsey JF (2008) Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning. J Glob Optim 42(4):587–607

Bevilacqua V, Mastronardi G, Piscopo G (2007) Evolutionary approach to inverse planning in coplanar radiotherapy. Image Vis Comput 25(2):196–203

Bortfeld T, Schlegel W (1993) Optimization of beam orientations in radiation therapy: some theoretical considerations. Phys Med Biol 38(2):291–304

Cheong K, Suh T, Romeijn H, Li J, Dempsey J (2005) Fast nonlinear optimization with simple bounds for IMRT planning. Med Phys 32:1975–1975

Craft D (2007) Local beam angle optimization with linear programming and gradient search. Phys Med Biol 52(7):N127–N135

Craft DL, Halabi TF, Shih HA, Bortfeld TR (2006) Approximating convex Pareto surfaces in multiobjective radiotherapy planning. Med Phys 33:3399–3407

D'Souza W, Meyer RR, Shi L (2004) Selection of beam orientations in intensity-modulated radiation therapy using single-beam indices and integer programming. Phys Med Biol 49(15):3465–3481

Das SK, Marks LB (1997) Selection of coplanar or noncoplanar beams using three-dimensional optimization based on maximum beam separation and minimized nontarget irradiation. Int J Radiat Oncol Biol Phys 38(3):643

Deasy J, Lee EK, Bortfeld T, Langer M, Zakarian K, Alaly J, Zhang Y, Liu H, Mohan R, Ahuja R (2006) A collaboratory for radiation therapy treatment planning optimization research. Ann Oper Res 148(1): 55–63

Deasy JO, Blanco AI, Clark VH (2003) CERR: a computational environment for radiotherapy research. Med Phys 30:979–985

Djajaputra D, Wu Q, Wu Y, Mohan R (2003) Algorithm and performance of a clinical IMRT beam-angle optimization system. Phys Med Biol 48(19):3191–3212

Ehrgott M, Holder A, Reese J (2008) "Beam selection in radiotherapy design". Linear Algebra Appl 428:1272–1312

Ehrgott M, Johnston R (2003) Optimisation of beam directions in intensity modulated radiation therapy planning. OR Spectr 25(2):251–264

El-Beltagy MA, Keane AJ (1999) Metalmodeling techniques for evolutionary optimization of computationally expensive problems: promises and limitations. In: Proceedings of the genetic and evolutionary computation conference, Orlando, USA, pp 196–203

Emmerich M, Giotis A, Özdemir M, Bäck T, Giannakoglou K (2002) Metamodel-assisted evolution strategies. Parallel problem solving from nature PPSN VII. Springer, Berlin

Fiege J, McCurdy B, Potrebko P, Champion H, Cull A (2011) PARETO: a novel evolutionary optimization approach to multiobjective IMRT planning. Med Phys 38:5217–5229

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

Goitein M, Abrams M, Rowell D, Pollari H, Wiles J (1983) Multi-dimensional treatment planning: II. Beam's eye-view, back projection, and projection through CT sections. Int J Radiat Oncol Biol Phys 9(6):789–797

Goodband JH, Haas OCL (2008). Artificial neural networks in radiation therapy. Intell Adapt Syst Med 213–257

Gulliford SL, Webb S, Rowbottom CG, Corne DW, Dearnaley DP (2004) Use of artificial neural networks to predict biological outcomes for patients receiving radical radiotherapy of the prostate. Radiother Oncol 71(1):3–12

Haas OCL, Burnham KJ, Mills JA (1998) Optimization of beam orientation in radiotherapy using planar geometry. Phys Med Biol 43:2179–2193

Haas OCL, Reeves CR (2005) Genetic algorithms in radiotherapy. Stud Multidiscip 3:447–482

Han J, Kamber M (2006) Data mining: concepts and techniques. Morgan Kaufmann, Burlington

Holdsworth C, Kim M, Liao J, Phillips MH (2010) A hierarchical evolutionary algorithm for multiobjective optimization in IMRT. Med Phys 37:4986–4997

Hüsken M, Jin Y, Sendhoff B (2005) Structure optimization of neural networks for evolutionary design optimization. Soft Comput A Fusion Found Methodol Appl 9(1):21–28

Jin Y, Branke J (2005) Evolutionary optimization in uncertain environments-a survey. Evol Comput IEEE Trans 9(3):303–317

Jin Y, Olhofer M, Sendhoff B (2002) A framework for evolutionary optimization with approximate fitness functions. Evol Comput IEEE Trans 6(5):481–494

Jin Y, Sendhoff B (2004) Reducing fitness evaluations using clustering techniques and neural network ensembles. Genetic and Evolutionary Computation GECCO. Springer, Seattle

Kalantzis G, Vasquez-Quino LA, Zalman T, Pratx G, Lei Y (2011) Toward IMRT 2D dose modeling using artificial neural networks: A feasibility study. Med Phys 38:5807–5817

Knezevic A (2008) Overlapping confidence intervals and statistical significance. StatNews #73. http://www.cscu.cornell.edu/news/statnews/stnews73.pdf. Accessed 14th January 2013

Knowles J, Gorne D, Bishop M (1998) Evolutionary training of artificial neural networks for radiotherapy treatment of cancers. IEEE World Congress on Computational Intelligence, pp 398–403, Alasca

Lahanas M, Schreibmann E, Milickovic N, Baltas D (2003) Intensity modulated beam radiation therapy dose optimization with multiobjective evolutionary algorithms. Evolutionary multi-criterion optimization, second international conference, pp 70–70, Springer, Faro, Portugal

Lee EK, Fox T, Crocker I (2003) Integer programming applied to intensity-modulated radiation therapy treatment planning. Ann Oper Res 119(1):165–181

Lee EK, Fox T, Crocker I (2006) Simultaneous beam geometry and intensity map optimization in intensity-modulated radiation therapy. Int J Radiat Oncol Biol Phys 64(1):301–320

Lei J, Li Y (2008). A DNA genetic algorithm for beam angle selection in radiotherapy planning. Cybernetics and Intelligent Systems, 2008 IEEE conference on, IEEE, pp 1331–1336

Li Y, Lei J (2010) A feasible solution to the beam-angle-optimization problem in radiotherapy planning with a DNA-based genetic algorithm. Biomed Eng IEEE Trans 57(3):499–508

Li Y, Yao D (2006) Accelerating the radiotherapy planning with a hybrid method of genetic algorithm and ant colony system. Lect Notes Comput Sci 4222:340

Li Y, Yao D, Yao J, Chen W (2005) A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning. Phys Med Biol 50:3491

Li Y, Yao D, Zheng J, Yao J (2006) A modified genetic algorithm for the beam angle optimization problem in intensity-modulated radiotherapy planning. Artif Evol 97–106

Li Y, Yao J, Yao D (2004) Automatic beam angle selection in IMRT planning using genetic algorithm. Phys Med Biol 49:1915

Lim GJ (2008). Introduction to radiation therapy planning optimization. Optim Med Biol 197–221

Lim GJ, Cao W (2012) A two-phase method for selecting IMRT treatment beam angles: Branch-and-Prune and local neighborhood search. Eur J Oper Res 217(3):609–618

Liu HH, Jauregui M, Zhang X, Wang X, Dong L, Mohan R (2006) Beam angle optimization and reduction for intensity-modulated radiation therapy of non-small-cell lung cancers. Int J Radiat Oncol Biol Phys 65(2):561

Llacer J, Li S, Agazaryan N, Promberger C, Solberg TD (2009) Non-coplanar automatic beam orientation selection in cranial IMRT: a practical methodology. Phys Med Biol 54:1337–1368

Lu HM, Kooy HM, Leber ZH, Ledoux RJ (1997) Optimized beam planning for linear accelerator-based stereotactic radiosurgery. Int J Radiat Oncol Biol, Phys 39(5):1183–1189

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X

Mathieu R, Martin E, Gschwind R, Makovicka L, Contassot-Vivier S, Bahi J (2005) Calculations of dose distributions using a neural network model. Phys Med Biol 50:1019–1028

Nazareth DP, Brunner S, Jones MD, Malhotra HK, Bakhtiari M (2009) Optimization of beam angles for intensity modulated radiation therapy treatment planning using genetic algorithm on a distributed computing platform. J Med Phys Assoc Med Phys India 34(3):129

Pugachev A, Xing L (2001a) Computer-assisted selection of coplanar beam orientations in intensity-modulated radiation therapy. Phys Med Biol 46(9):2467–2476

Pugachev A, Xing L (2001b) Pseudo beam's eye-view as applied to beam orientation selection in intensity-modulated radiation therapy. Int J Radiat Oncol Biol Phys 51(5):1361–1370

Pugachev A, Xing L (2002) Incorporating prior knowledge into beam orientaton optimization in IMRT. Int J Radiat Oncol Biol Phys 54(5):1565–1574

Rocha H, Dias J, Ferreira BC, Lopes MC (2012) Selection of intensity modulated radiation therapy treatment beam directions using radial basis functions within a pattern search methods framework. J Glob Optim (accepted)

Romeijn HE, Ahuja RK, Dempsey JF, Kumar A, Li JG (2003) A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. Phys Med Biol 48(21):3521–3542

Rowbottom CG, Webb S, Oldham M (1999) Beam-orientation customization using an artificial neural network. Phys Med Biol 44:2251–2262

Schreibmann E, Lahanas M, Xing L, Baltas D (2004) Multiobjective evolutionary optimization of the number of beams, their orientations and weights for intensity-modulated radiation therapy. Phys Med Biol 49(5):747–770

Shanker M, Hu MY, Hung MS (1996) Effect of data standardization on neural network training. Omega 24(4):385–397

Stein J, Mohan R, Wang XH, Bortfeld T, Wu Q, Preiser K, Ling CC, Schlegel W (1997) Number and orientations of beams in intensity-modulated radiation treatments. Med Phys 24:149–160

Vasseur A, Makovicka L, Martin Ã, Sauget M, Contassot-Vivier S, Bahi J (2008) Dose calculations using artificial neural networks: a feasibility study for photon beams. Nucl Instrum Methods Phys Res Sect B Beam Interact Mater Atoms 266(7):1085–1093

Wells DM, Niederer J (1998) A medical expert system approach using artificial neural networks for standardized treatment planning. Int J Radiat Oncol Biol Phys 41(1):173–182

Willoughby TR, Starkschall G, Janjan NA, Rosen II (1996) Evaluation and scoring of radiotherapy treatment plans using an artificial neural network. Int J Radiat Oncol Biol Phys 34(4):923–930

Witten IH, Frank E (2005) Data mining: practical machine learning tools and techniques. Morgan Kaufmann, Burlington

Wu X, Zhu Y, Dai J, Wang Z (2000) Selection and determination of beam weights based on genetic algorithms for conformal radiotherapy treatment planning. Phys Med Biol 45:2547–2558

Journal: 10100 Article No.: 0289 TYPESET DISK LE CP Disp.:2013/3/19 Pages: 25 Layout: Small-X