Incorporating Radial Basis Functions in Pattern Search Methods: Application to Beam Angle Optimization in Radiotherapy Treatment Planning

Humberto Rocha[1], Joana M. Dias[1,2], Brigida C. Ferreira[3,4], and Maria do Carmo Lopes[3,4]

[1] INESC-Coimbra, Rua Antero de Quental, 199
3000-033 Coimbra, Portugal
[2] Faculdade de Economia, Universidade de Coimbra,
3004-512 Coimbra, Portugal
[3] I3N, Departamento de Física, Universidade de Aveiro,
3810-193 Aveiro, Portugal
[4] Serviço de Física Médica, IPOC-FG, EPE,
3000-075 Coimbra, Portugal
hrocha@mat.uc.pt, joana@fe.uc.pt, brigida@ua.pt,
mclopes@ipocoimbra.min-saude.pt

Abstract. The global optimization of black-box functions with many local minima occurs in many branches of science and engineering. There are many methods and heuristics to address this type of problems. However, for problems with expensive black-box functions, both in terms of cost or time, the number of function evaluations required by most of the methods or heuristics is prohibitive. The pattern search methods framework is suited to address this type of problems since it requires few function value evaluations to converge and have the ability to avoid local entrapment. The ability of this class of methods to obtain global minima depends on the incorporation of methods or heuristics for global optimization on their, so called, search step. We propose the use of radial basis functions both to influence the quality of the local minimizer found by the method and also to obtain a better coverage of the search space. Our approach is tailored for addressing the beam angle optimization (BAO) problem in intensity modulated radiation therapy treatment planning, but can be easily extended for other general problems. The BAO problem is quite difficult, and yet to be solved in a satisfactory way, since it is a highly non-convex optimization problem with many local minima. A couple of retrospective treated cases of head-and-neck tumors at the Portuguese Institute of Oncology of Coimbra is used to discuss the benefits of using our approach in the optimization of the BAO problem.

Keywords: Pattern Search Methods, Radial Basis Functions, Radiotherapy, IMRT, Beam Angle Optimization.

B. Murgante et al. (Eds.): ICCSA 2012, Part III, LNCS 7335, pp. 1–16, 2012.
© Springer-Verlag Berlin Heidelberg 2012

1 Introduction

The global optimization of black-box functions with many local minima occurs in many branches of science and engineering. Directional direct-search methods have been used to tackle this type of problems [1]. The pattern search methods framework is the most used and implemented class of directional direct-search methods [3]. Pattern search methods are organized around two steps at every iteration: the poll step and the search step. The poll step performs a local search in a neighborhood around the current iterate using the concepts of positive bases, and under the appropriate assumptions, it guarantees global convergence to stationary points. The search step consists of a finite search, free of rules, away from the current iterate, and the ability to obtain global minima depends on the incorporation of methods or heuristics for global optimization. An example of such hybridization is the use of particle swarm optimization in the search step of the pattern search methods framework [19]. However, for problems with expensive black-box functions, both in terms of cost or time, the number of function evaluations required for this type of strategies is prohibitive for obtaining results in an acceptable time frame or within the budget. The beam angle optimization (BAO) problem in intensity modulated radiation therapy treatment planning is such problem and will be used to illustrate the merits of our approach. The intensity modulated radiation therapy (IMRT) is a modern type of radiation therapy, whose inverse planning leads to complex optimization problems, including the BAO problem - the problem of deciding which incidence radiation beam angles should be used. The BAO problem is quite difficult, and yet to be solved in a satisfactory way, since it is a highly non-convex optimization problem with many local minima [5]. Therefore, methods that avoid being easily trapped in local minima should be used. Moreover, each function evaluation is time expensive so methods that require few function value evaluations should be used to tackle the BAO problem. The pattern search methods framework is suited to address the BAO problem since it requires few function value evaluations to converge and have the ability to avoid local entrapment. Here, we will discuss the benefits of incorporating radial basis functions in the pattern search methods framework for the optimization of the highly non-convex BAO problem. Radial basis functions are used both to influence the quality of the local minimizer found by the method and also to obtain a better coverage of the search space in amplitude. A couple of retrospective treated cases of head-and-neck tumors at the Portuguese Institute of Oncology of Coimbra is used to discuss the benefits of using our approach in the optimization of the BAO problem. Our approach is tailored to address this particular problem but it can be easily extended for other general problems. The paper is organized as follows. In the next section we describe the BAO problem. Radial basis functions interpolation and its use within the pattern search methods framework is presented in section 3. Clinical examples of head-and-neck cases used in the computational tests are presented in section 4. Section 5 presents the experimental results. In the last section we have the conclusions.

2 Beam Angle Optimization in IMRT Treatment Planning

The purpose of radiation therapy is to deliver a dose of radiation to the tumor volume to sterilize all cancer cells minimizing the collateral effects on the surrounding healthy organs and tissues. Typically, radiation is generated by a linear accelerator mounted on a gantry that can rotate along a central axis and is delivered with the patient immobilized on a couch that can rotate. The rotation of the couch combined with the rotation of the gantry allows radiation from almost any angle around the tumor. In IMRT the radiation beam is modulated by a multileaf collimator that enables the transformation of the beam into a grid of smaller beamlets of independent intensities. A common way to solve the inverse planning in IMRT optimization problems is to use a beamlet-based approach leading to a large-scale programming problem. Due to the complexity of the whole optimization problem, many times the treatment planning is divided into three smaller problems which can be solved sequentially: BAO problem, fluence map optimization (FMO) problem, and leaf sequencing problem. Here, we will focus our attention in the BAO problem, using coplanar angles, and we will assume that the number of beam angles is defined a priori by the treatment planner.

Many attempts to address the BAO problem can be found in the literature including simulated annealing [4], genetic algorithms [10], particle swarm optimization [12] or other heuristics incorporating a priori knowledge of the problem. Although those global heuristics can theoretically avoid local optima, globally optimal or even clinically better solutions can not be obtained without a large number of objective function evaluations. For that reason, many of the previous BAO studies are based on a variety of scoring methods or approximations to the FMO to gauge the quality of the beam angle set. 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. Therefore, our approach for modeling the BAO problem, similarly to [2,5], uses the optimal solution value of the FMO problem as the measure of the quality for a given beam angle set. Thus, we will present the formulation of the BAO problem followed by the formulation of the FMO problem we used.

2.1 BAO Model

Let us consider n to be the fixed number of (coplanar) beam directions, i.e., n beam angles are chosen on a circle around the CT-slice of the body that contains the isocenter (usually the center of mass of the tumor). Typically, the BAO problem is formulated as a combinatorial optimization problem in which a specified number of beam angles is to be selected among a beam angle candidate pool. The continuous $[0°, 360°]$ gantry angles are generally discretized into equally spaced directions with a given angle increment, such as 5 or 10 degrees. We will consider a different approach for the formulation of the BAO problem.

All continuous $[0°, 360°]$ gantry angles will be considered instead of a discretized sample. Since the angle $-5°$ is equivalent to the angle $355°$ and the angle $365°$ is the same as the angle $5°$, we can avoid a bounded formulation. A basic formulation for the BAO problem is obtained by selecting an objective function such that the best set of beam angles is obtained for the function's minimum:

$$\min f(\theta_1, \ldots, \theta_n)$$

$$s.t. \ (\theta_1, \ldots, \theta_n) \in \mathbb{R}^n.$$

Here, the objective $f(\theta_1, \ldots, \theta_n)$ that measures the quality of the set of beam directions $\theta_1, \ldots, \theta_n$ is the optimal value of the FMO problem for each fixed set of beam directions. Such functions have numerous local optima, which increases the difficulty of obtaining a good global solution. Thus, the choice of the solution method becomes a critical aspect for obtaining a good solution. Our formulation was mainly motivated by the ability of using a class of solution methods that we consider to be suited to successfully address the BAO problem: pattern search methods. The FMO model used is presented next.

2.2 FMO Model

For a given beam angle set, an optimal IMRT plan is obtained by solving the FMO problem - the problem of determining the optimal beamlet weights for the fixed beam angles. Many mathematical optimization models and algorithms have been proposed for the FMO problem, including linear models [17], mixed integer linear models [11] and nonlinear models [2].

Radiation dose distribution deposited in the patient, measured in Gray (Gy), needs to be assessed accurately in order to solve the FMO problem, i.e., to determine optimal fluence maps. Each structure's volume is discretized into 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 constructed from the collection of all beamlet weights, 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 (N_v) and the number of columns equals the number of beamlets (N_b) from all beam directions considered. Therefore, using matrix format, we can say that the total dose received by the voxel i is given by $\sum_{j=1}^{N_b} D_{ij} w_j$, with w_j the weight 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.

Here, we will use a convex penalty function voxel-based nonlinear model [2]. In this model, each voxel is penalized according to 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 non-negativity constraints on the fluence values [17]:

$$\min_w \sum_{i=1}^{N_v} \frac{1}{v_S} \left[\underline{\lambda}_i \left(T_i - \sum_{j=1}^{N_b} D_{ij} w_j \right)_+^2 + \overline{\lambda}_i \left(\sum_{j=1}^{N_b} D_{ij} w_j - T_i \right)_+^2 \right]$$

$$s.t. \quad w_j \geq 0,\ j = 1, \ldots, N_b,$$

where T_i is the desired dose for voxel i, $\underline{\lambda}_i$ and $\overline{\lambda}_i$ are the penalty weights of underdose and overdose of voxel i, and $(\cdot)_+ = \max\{0, \cdot\}$. Although this formulation allows unique weights for each voxel, similarly to the implementation in [2], 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 (v_S). 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 [2].

The FMO model is used as a black-box function. It is beyond the scope of this study to discuss if this formulation of the FMO problem is preferable to others. The conclusions drawn regarding BAO coupled with this nonlinear model are valid also if different FMO formulations are considered.

3 Radial Basis Function Interpolation and Its Use within the Pattern Search Methods Framework

For numerical approximation of multivariate functions, radial basis functions (RBFs) can provide excellent interpolants. For any finite data set in any Euclidean space, one can construct an interpolation of the data by using RBFs, even if the data points are unevenly and sporadically distributed in a high dimensional Euclidean space. However, RBF interpolant trends between and beyond the data points depend on the RBF used and may exhibit undesirable trends using some RBFs while the trends may be desirable using other RBFs. Numerical choice of the most adequate RBF for the problem at hand should be done instead of an usual a priori choice [15]. Next, we will formulate RBF interpolation problems and describe the strategy used to take advantage of the incorporation of RBF interpolants in the pattern search method framework applied to the BAO problem.

3.1 RBF Interpolation Problems

Let $f(\mathbf{x})$ be the true response to a given input vector $\mathbf{x}$ (of n components) such that the value of f is only known at a set of N input vectors $\mathbf{x} = \mathbf{x}^1, \ldots, \mathbf{x}^N$, i.e., only $f(\mathbf{x}^k)$ $(k = 1, \ldots, N)$ are known. An interpolation model $g(\mathbf{x})$ generated from a RBF $\varphi(t)$ can be represented in the following form:

$$g(\mathbf{x}) = \sum_{j=1}^{N} \alpha_j \varphi(\|\mathbf{x} - \mathbf{x}^j\|), \tag{1}$$

where α_j are the coefficients to be determined by interpolation conditions, $g(\mathbf{x}^k) = f(\mathbf{x}^k)$ $(k = 1, \ldots, N)$, $\|\mathbf{x} - \mathbf{x}^j\|$ denotes the parameterized distance between $\mathbf{x}$ and $\mathbf{x}^j$ defined as $||\mathbf{x} - \mathbf{x}^j|| = \sqrt{\sum_{i=1}^n |\theta_i|\left(x_i - x_i^j\right)^2}$, and $\theta_1, \ldots, \theta_n$ are scalars [15]. For fixed parameters θ_i, the coefficients $\alpha_1, \ldots, \alpha_N$ in Eq. (1) can be calculated by solving the following linear system of interpolation equations:

$$\sum_{j=1}^{N} \alpha_j \varphi(||\mathbf{x}^k - \mathbf{x}^j||) = f(\mathbf{x}^k), \quad \text{for } k = 1, \ldots, N. \tag{2}$$

The most popular examples of RBF [14] are cubic spline $\varphi(t) = t^3$, thin plate spline $\varphi(t) = t^2 \ln t$, multiquadric $\varphi(t) = \sqrt{1 + t^2}$, and Gaussian $\varphi(t) = \exp(-\mathrm{t}^2)$. These RBFs can be used to model cubic, almost quadratic, and linear growth rates, as well as exponential decay, of the response for trend predictions. A unique interpolant is guaranteed for multiquadric and Gaussian RBFs, (i.e., the system matrix in Eq. (2) is nonsingular) even if the input vectors $\mathbf{x}^j$ are few and poorly distributed, provided only that the input vectors are all different when $N > 1$. However, for cubic and thin plate spline RBFs, the system matrix in Eq. (2) might be singular [14]. An easy way to avoid this problem on the cubic and thin plate spline RBF interpolants is to add low-degree polynomials to interpolation functions in Eq. (1) (see [15]).

The constructed interpolant $g(\mathbf{x})$ in Eq. (1) depends on "subjective" choice of $\varphi(t)$, and model parameters $\theta_1, \ldots, \theta_n$. While one can try all the possible choices of $\varphi(t)$ in search of a desirable interpolant, there are infinitely many choices for $\theta_1, \ldots, \theta_n$. Mathematically, one could pick any fixed set of $\theta_1, \ldots, \theta_n$ and construct the interpolation function for the given data. However, two different sets of $\theta_1, \ldots, \theta_n$ will lead to two interpolation models that behave very differently between the input vectors $\mathbf{x}^1, \ldots, \mathbf{x}^N$. Model parameter tuning for RBF interpolation aims at finding a set of parameters $\theta_1, \ldots, \theta_n$ that results in the best prediction model of the unknown response based on the available data. The prediction accuracy can be used as a criterion for choosing the best basis function $\varphi(t)$ and parameters θ_i. Cross-validation (CV) [18] was proposed to find $\varphi(t)$ and θ_i that lead to an approximate response model $g(\mathbf{x})$ with optimal prediction capability and proved to be effective [18]. The leave-one-out CV procedure is usually used in model parameter tuning for RBF interpolation [18]:

Algorithm 1. (Leave-one-out cross-validation for RBF interpolation)

1. Fix a set of parameters $\theta_1, \ldots, \theta_n$.
2. For $j = 1, \ldots, N$, construct the RBF interpolant $g_{-j}(\mathbf{x})$ of the data points $(\mathbf{x}^k, f(\mathbf{x}^k))$ for $1 \leq k \leq N, k \neq j$.
3. Use the following CV root mean square error as the prediction error:

$$E^{CV}(\theta_1, \ldots, \theta_n) = \sqrt{\frac{1}{N} \sum_{j=1}^{N} (g_{-j}(\mathbf{x}^j) - f(\mathbf{x}^j))^2}. \tag{3}$$

The goal of model parameter tuning by CV is to find $\theta_1, \ldots, \theta_n$ that minimize the CV error, $E^{CV}(\theta_1, \ldots, \theta_n)$, so that the interpolation model has the highest prediction accuracy when CV error is the measure. Using different θ_i allows the model parameter tuning to scale each variable x_i based on its significance in modeling the variance in the response, thus, has the benefit of implicit variable screening built in the model parameter tuning.

3.2 Incorporation of RBF Models in the Pattern Search Methods Framework Tailored for the BAO Problem

Pattern search methods are directional direct search methods that belong to a broader class of derivative-free optimization methods, such that iterate progression is solely based on a finite number of function evaluations in each iteration, without explicit or implicit use of derivatives. Pattern search methods generate a sequence of non-increasing iterates $\{\mathbf{x}^k\}$ using positive bases (or positive spanning sets) and moving towards a direction that would produce a function decrease. A positive basis for $\mathbb{R}^n$ can be defined as a set of nonzero vectors of $\mathbb{R}^n$ whose positive combinations span $\mathbb{R}^n$ (positive spanning set), but no proper set does. A positive spanning set contains at least one positive basis. It can be shown that a positive basis for $\mathbb{R}^n$ contains at least $n+1$ vectors and cannot contain more than $2n$ [8]. Positive basis with $n+1$ and $2n$ elements are referred to as minimal and maximal positive basis, respectively. Commonly used minimal and maximal positive basis are $[I \;\; -e]$, with I being the identity matrix of dimension n and $e = [1 \; \ldots \; 1]^\top$, and $[I \;\; -I]$, respectively.

One of the main features of positive bases (or positive spanning sets), that is the motivation for directional direct search methods, is that, unless the current iterate is at a stationary point, there is always a vector $\mathbf{v}^i$ in a positive basis (or positive spanning set) that is a descent direction [8], i.e., there is an $\alpha > 0$ such that $f(\mathbf{x}^k + \alpha \mathbf{v}^i) < f(\mathbf{x}^k)$. This is the core of directional direct search methods and in particular of pattern search methods. The notions and motivations for the use of positive bases, its properties and examples can be found in [1,8].

Pattern search methods framework is briefly presented next. Let us denote by $\mathbf{V}$ the $n \times p$ matrix whose columns correspond to the p ($\geq n+1$) vectors forming a positive spanning set. Given the current iterate $\mathbf{x}^k$, at each iteration k, the next point $\mathbf{x}^{k+1}$, aiming to provide a decrease of the objective function, is chosen from a finite number of candidates on a given mesh $M_k = \{\mathbf{x}^k + \alpha_k \mathbf{V}\mathbf{z} : \mathbf{z} \in \mathbb{Z}_+^p\}$, where α_k is the mesh-size (or step-size) parameter and $\mathbb{Z}_+$ is the set of nonnegative integers. Pattern search methods are organized around two steps at every iteration. The first step consists of a finite search on the mesh, free of rules, with the goal of finding a new iterate that decreases the value of the objective function at the current iterate. This step, called the search step, has the flexibility to use any strategy, method or heuristic, or take advantage of a priori knowledge of the problem at hand, as long as it searches only a finite number of points in the mesh. The search step provides the flexibility for a global search since it allows searches away from the neighborhood of the current iterate, and influences the quality of the local minimizer or stationary point found by the method. If the search step fails to produce a decrease in the objective function, a second step, called the poll step, is performed around the current iterate. The poll step follows stricter rules and, using the concepts of positive bases, attempts to perform a local search in a mesh neighborhood around $\mathbf{x}^k$, $\mathcal{N}(\mathbf{x}^k) = \{\mathbf{x}^k + \alpha_k \mathbf{v} : \text{ for all } \mathbf{v} \in P_k\} \subset M_k$, where P_k is a positive basis chosen from the finite positive spanning set $\mathbf{V}$. For a sufficiently small mesh-size parameter α_k, the poll step is guaranteed to provide a function reduction, unless the current iterate is at a stationary point [1]. So, if the poll step also fails to produce a function reduction, the mesh-size parameter α_k must be decreased. On the other hand, if both the search and poll steps fail to obtain an improved value for the objective function, the mesh-size parameter is increased or held constant.

The most common choice for the mesh-size parameter update is to half the mesh-size parameter at unsuccessful iterations and to keep it or double it at successful ones. Note that, if the initial mesh parameter is a power of 2, ($\alpha_0 = 2^l, l \in \mathbb{N}$), and the initial point is a vector of integers, using this common mesh update, all iterates will be a vector of integers until the mesh-size parameter becomes inferior to 1. This possibility is rather interesting for the BAO problem.

Recently, the efficiency of pattern search methods improved significantly by reordering the poll directions according to descent indicators built from simplex gradients [7]. Here, the poll directions are reordered according to the RBF model values. The most common approach for incorporating interpolation models in the search step consists of forming an interpolation model and finding its minimum. For example, in Custódio et al. [6], the search step computes a single trial point using minimum Frobenius norm quadratic models to be minimized within a trust region. The size of the trust region is coupled to the radius of the sample set. Thus, for an effective global search, the sample points should span all the search space. That could be achieved by using larger initial step-size parameters. However, since the BAO problem has many local minima and the number of sample points is scarce, the polynomial interpolation or regression models (usually quadratic models) used within the trust region struggle to find the best local minima. Therefore, starting with larger mesh-size parameters may lead to similar or worst results obtained when starting with smaller mesh-size parameters and at the cost of more function value evaluations [16]. An alternative and popular approach to keep small mesh-size parameters and still have a good coverage of the whole search space is to use a multi-start approach. However, the multi-start approach has the disadvantage of increasing the total number of function evaluations and with that the overall computational time. Moreover, the obtained good span of $\mathbb{R}^2$ in amplitude is only obtained by overlapping all the iterates giving the illusion that unusual beam angle configurations were tested while in fact only local searches around the initial beam angle configurations were performed. We adopted a different strategy, by considering a single starting point, a small initial mesh-size parameter, and trying to obtain a good span in amplitude of $\mathbb{R}^2$ by incorporating radial basis functions models in the search step. The strategy sketched here is tailored for addressing the BAO problem and does not include the formal minimization of the RBF model:

Algorithm 2. (PSM framework using RBFs for the BAO problem)

0. Initialization Set $k = 0$. Choose $\mathbf{x}^0 \in \mathbb{R}^n$, $\alpha_0 > 0$, and a positive spanning set $\mathbf{V}$.

1. Search step If the number of evaluated points is not greater than $n+1$ skip the search step. Otherwise, build a RBF model and while a decrease on the objective function value is not achieved, compute the RBFs trial points:
For each beam angle direction ($i = 1, \ldots, n$)

- **a.** Evaluate the RBF model for every degree between the previous beam direction and the next one.
- **b.** Find the minimum of those values that correspond to a beam direction that was not evaluated yet and, is at least 4 degrees away from a previously evaluated one, for the beam direction at stake.
- **c.** Take as RBF trial point the current iterate updating the beam direction corresponding to the minimum found in b.

If no RBF trial point correspond to a decrease on the objective function value, go to step 2 and the search step is declared unsuccessful. Otherwise, go to step 4 and both the search step and iteration are declared successful.

2. Poll step This step is only performed if the search step is unsuccessful. If a RBF model was computed in the previous step then reorder the poll directions according to the RBF model values. If $f(\mathbf{x}^k) \leq f(\mathbf{x})$ for every $\mathbf{x}$ in the mesh neighborhood $\mathcal{N}(\mathbf{x}^k)$, then go to step 3 and shrink M_k. Both poll step and iteration are declared unsuccessful. Otherwise, choose a point $\mathbf{x}^{k+1} \in \mathcal{N}(\mathbf{x}^k)$ such that $f(\mathbf{x}^{k+1}) < f(\mathbf{x}^k)$ and go to step 4. Both poll step and iteration are declared successful.

3. Mesh reduction Let $\alpha_{k+1} = \frac{1}{2} \times \alpha_k$. Set $k = k + 1$ and return to step 1.

4. Mesh expansion Let $\alpha_{k+1} = \alpha_k$. Set $k = k + 1$ and return to step 1.

Our main goal for using a RBF model in the search step of the pattern search methods framework is to properly explore the search space in amplitude without a random criteria. Therefore, each beam direction is tested every degree between the previous beam direction and the next one as stated in step a of the RBF trial points computation. A proper minimization is unnecessary since we are interested in integer beam angle directions. Step b of the RBF trial points computation within search step forces the algorithm to consider only directions in regions not yet explored which is the main goal of the RBF models here (directions that are less than 4 degrees apart are considered to be clinically equivalent). The maximum number of points computed in the search step is n (e.g. if the search step is unsuccessful). More conservative strategies could be adopted considering, e.g., only the best of the RBF trial points.

The benefits of using RBFs in the pattern search methods framework for the optimization of the BAO problem are illustrated using a set of clinical examples of head-and-neck cases that are presented next.

4 Head-and-Neck Clinical Examples

Two clinical examples of retrospective treated cases of head-and-neck tumors at the Portuguese Institute of Oncology of Coimbra (IPOC) are used to test the incorporation of RBF models in a pattern search methods framework. 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 with the typical 7-beam equispaced coplanar treatment plans. The patients' CT sets and delineated structures were exported via Dicom RT to a freeware computational environment for radiotherapy research (see Figure 1). 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 type 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 optimization.

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.

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 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.

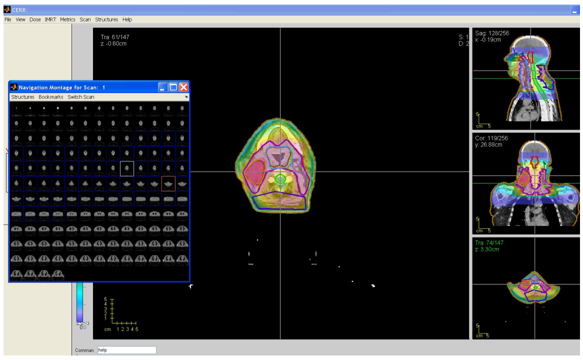

Fig. 1. Illustration of the structures visualized in CERR

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

Structure	Mean dose	Max 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	–

5 Results

Our tests were performed on a 2.66Ghz Intel Core Duo PC with 3 GB RAM. 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 [9] are used widely for IMRT treatment planning research. We used CERR 3.2.2 version and MATLAB 7.4.0 (R2007a). The dose was computed using CERR's pencil beam algorithm (QIB). 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 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.

We choose to implement the use of RBFs taking advantage of the availability of an existing pattern search methods framework implementation used successfully by us to tackle the BAO problem [16] – the last version of SID-PSM [6,7].

The spanning set used was the positive spanning set ($[e \ -e \ I \ -I]$, with I being the identity matrix and $e = [1 \ \ldots \ 1]^T$). Each of these directions corresponds to, respectively, the rotation of all incidence directions clockwise, the rotation of all incidence directions counter-clockwise, the rotation of each individual incidence direction clockwise, and the rotation of each individual incidence direction counter-clockwise. The initial mesh-size parameter was set to $\alpha_0 = 4$ since larger values increase the number of function evaluations with no benefits [16]. Since the initial points were integer vectors, all iterates will have integer values as long as the mesh parameter does not become less than one. Therefore, the stopping criteria adopted was the mesh parameter becoming less than one.

The RBFs incorporation into the pattern search methods framework was tested using two 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 five to nine equally spaced coplanar orientations around the patient. Treatment plans with seven equispaced coplanar beams were used at IPOC and are commonly used in practice to treat head-and-neck cases [2]. Therefore, treatment plans of seven coplanar orientations were obtained using our BAO algorithms, denoted *SID-PSM* and *PSM-RBF*, wether the algorithm used was the pattern search framework alone or incorporating RBFs, respectively. These treatment plans were compared with the typical 7-beam equispaced coplanar treatment plans denoted *equi*.

The main goal of the present work is to verify the contribution of the incorporation of RBF models in pattern search methods applied to the optimization of the BAO problem, both in terms of optimal function value found and appropriate search space coverage. Beforehand, we need to decide which RBF is better and should be used for the BAO problem. The CV error of an interpolation model can be a useful and objective tool to decide which RBF model is better. We used the MATLAB code *fminsearch*, an implementation of the Nelder-Mead [13] multidimensional search algorithm, to minimize the CV error $E^{CV}(\theta_1, \ldots, \theta_n)$ in Eq. (3) and to find the best model parameters $\theta_1, \ldots, \theta_n$. Instead of choosing a priori which RBF should be used, the RBF model used at each iteration is the one that yields the smallest CV error, and consequently the RBF model with the highest prediction accuracy.

The objective function value decrease versus the number of function evaluations required is presented in Fig. 2 to compare the performances of *SID-PSM* and *PSM-RBF*. By simple inspection we conclude that *PSM-RBF* leads to better optimal objective function values compared to *SID-PSM*. The results are presented in terms of number of function evaluations instead of overall computational time since for different dose engines, beamlet optimization methods or even other objective function strategies, the overall computational time may have a totally different magnitude. Dose computation using QIB consumed most of the overall computational time. In average it took two and five hours to run the BAO optimization using the *SID-PSM* and the *PSM-RBF* algorithms, respectively. Our objective is to emphasize the small number of function evaluations required by

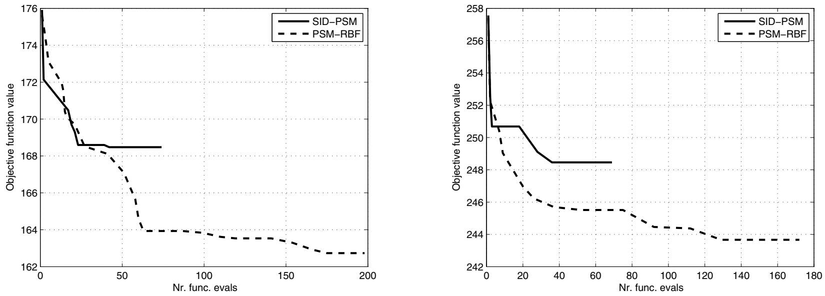

(a) (b)

Fig. 2. History of the 7-beam angle optimization process using *SID-PSM* and *PSM-RBF* for cases 1 and 2, 2(a) and 2(b) respectively

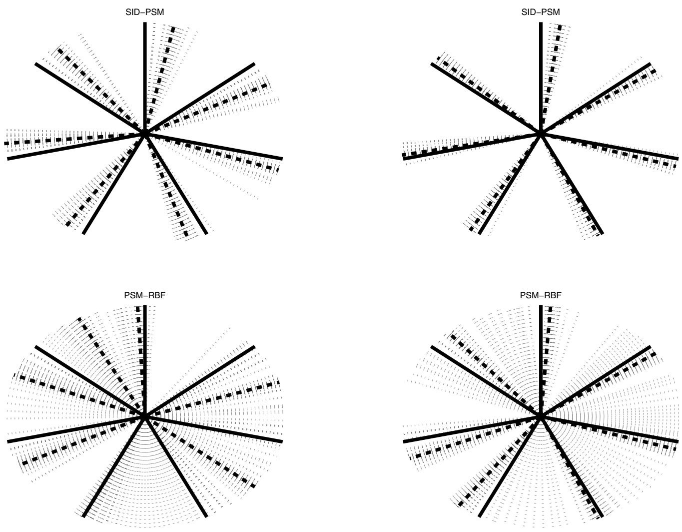

(a) (b)

Fig. 3. History of the 7-beam angle optimization process using *SID-PSM* and *PSM-RBF* for cases 1 and 2, 3(a) and 3(b) respectively. Initial angle configuration, optimal angle configuration and intermediate angle configurations are displayed with solid, dashed and dotted lines, respectively.

pattern search methods, compared to most of the global search methods, heuristics or strategies, even when using RBFs within the search step.

The history of the 7-beam angle optimization process using *SID-PSM* and *PSM-RBF*, in terms of beam directions tested, for each case, is presented in Fig. 3. By simple inspection we can verify that the sequence of iterates are better distributed by amplitude in $\mathbb{R}^2$ when using *PSM-RBF*, with a more appropriate coverage in amplitude of the whole search space.

Despite the improvement in FMO value, the quality of the results can be perceived considering a variety of metrics. 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. Another 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. DVH results for the two cases are displayed in Fig. 4. Since parotids are the most difficult organs to spare, and all the treatment plans fulfill the maximum dose requirements for the spinal cord and the brainstem, 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 doses. We can verify that all treatment plans obtained a satisfactory target coverage. However, as expected, the main differences reside in parotid sparing with clear advantage for the optimized treatment plans. In average, *SID-PSM* treatment plans reduced the parotid's mean dose irradiation in 0.8 Gy compared to the *equi* treatment

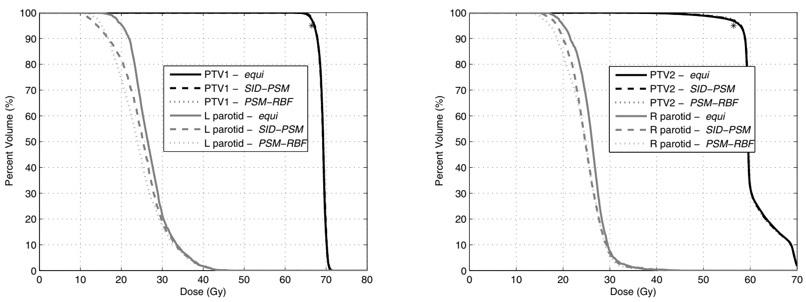

(a)

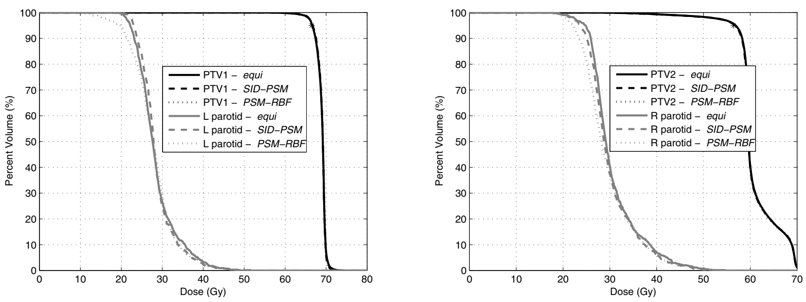

(b)

Fig. 4. Cumulative dose volume histogram comparing the results obtained by *equi*, *SID-PSM* and *PSM-RBF* for cases 1 and 2, 4(a) and 4(b) respectively

plans while *PSM-RBF* treatment plans reduced the parotid's mean dose irradiation in 1.5 Gy compared to the *equi* treatment plans. The differences between *SID-PSM* treatment plans and *PSM-RBF* treatment plans, concerning parotid sparing, show a clear advantage for the *PSM-RBF* treatment plans. The results displayed in Fig. 4 confirm the benefits of using the optimized beam directions, in particular using the directions obtained and used in *PSM-RBF* treatment plan.

6 Conclusions

The benefits of a tailored incorporation of RBFs in a pattern search methods framework were tested for the BAO problem using a couple of clinical head-and-neck cases. The BAO problem is a continuous global highly non-convex optimization problem known to be extremely challenging and yet to be solved satisfactorily. Pattern search methods are suited for the BAO problem since they require few function value evaluations and, similarly to other derivative-free optimization methods, have the ability to avoid local entrapment. The pattern search methods approach seems to be similar to neighborhood search approaches in which the neighborhood is constructed using the pattern search method. However, local neighborhood search approaches are only similar to the poll step of the pattern search methods framework. The existence of a search step with the flexibility to use any strategy, method or heuristic, or take advantage of a priori knowledge of the problem at hand, is an advantage that was explored successfully in this work. We have shown that a beam angle set can be locally improved in a continuous manner using pattern search methods. Moreover, it was shown that the incorporation of RBFs in the search step leads to an improvement of the local solution obtained. For numerical approximation of multivariate functions, RBFs can provide excellent interpolants, even if the data points available are unevenly and sporadically distributed. For the retrospective tumor cases tested, our RBFs tailored approach showed a positive influence on the quality of the local minimizer found and a clearly better coverage of the whole search space in amplitude. The improvement of the local solutions in terms of objective function value corresponded, for the head-and-neck cases tested, to high quality treatment plans with good target coverage and with improved organ sparing, in particular better parotid sparing. Moreover, we have to highlight the low number of function evaluations required to obtain locally optimal solutions, which is a major advantage compared to other global heuristics. This advantage should be even more relevant when considering non-coplanar directions since the number of possible directions to consider increase significantly. The efficiency on the number of function value computations is of the utmost importance for the optimization of other general expensive highly non-convex black-box functions.

Acknowledgements. This work was supported by FEDER funds through the COMPETE program and Portuguese funds through FCT under project grant PTDC/EIA-CCO/121450/2010 and by FCT under project grant PEst-C/EEI/UI0308/2011.

References

1. Alberto, P., Nogueira, F., Rocha, H., Vicente, L.N.: Pattern search methods for user-provided points: Application to molecular geometry problems. SIAM J. Optim. 14, 1216–1236 (2004)
2. 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. J. Global Optim. 42, 587–607 (2008)
3. Audet, C., Dennis Jr., J.E.: Analysis of generalized pattern search methods. SIAM J. Optim. 13, 889–903 (2003)
4. Bortfeld, T., Schlegel, W.: Optimization of beam orientations in radiation therapy: some theoretical considerations. Phys. Med. Biol. 38, 291–304 (1993)
5. Craft, D.: Local beam angle optimization with linear programming and gradient search. Phys. Med. Biol. 52, 127–135 (2007)
6. Custódio, A.L., Rocha, H., Vicente, L.N.: Incorporating minimum Frobenius norm models in direct search. Comput. Optim. Appl. 46, 265–278 (2010)
7. Custódio, A.L., Vicente, L.N.: Using sampling and simplex derivatives in pattern search methods. SIAM J. Optim. 18, 537–555 (2007)
8. Davis, C.: Theory of positive linear dependence. Am. J. Math. 76, 733–746 (1954)
9. Deasy, J.O., Blanco, A.I., Clark, V.H.: CERR: A Computational Environment for Radiotherapy Research. Med. Phys. 30, 979–985 (2003)
10. Ehrgott, M., Holder, A., Reese, J.: Beam selection in radiotherapy design. Linear Algebra Appl. 428, 1272–1312 (2008)
11. Lee, E.K., Fox, T., Crocker, I.: Integer programming applied to intensity-modulated radiation therapy treatment planning. Ann. Oper. Res. 119, 165–181 (2003)
12. Li, Y., Yao, D., Yao, J., Chen, W.: A particle swarm optimization algorithm for beam angle selection in intensity modulated radiotherapy planning. Phys. Med. Biol. 50, 3491–3514 (2005)
13. Nelder, J.A., Mead, R.: A simplex method for function minimization. Comput. J. 7, 308–313 (1965)
14. Powell, M.: Radial Basis Function Methods for Interpolation to Functions of Many Variables. HERMIS: Int. J. Computer Maths & Appl. 3, 1–23 (2002)
15. Rocha, H.: On the selection of the most adequate radial basis function. Appl. Math. Model. 33, 1573–1583 (2009)
16. Rocha, H., Dias, J.M., Ferreira, B.C., Lopes, M.C.: Beam angle optimization using pattern search methods: initial mesh-size considerations. In: Proceedings of the 1st International Conference on Operations Research and Enterprise Systems (2012)
17. Romeijn, H.E., Ahuja, R.K., Dempsey, J.F., Kumar, A., Li, J.: A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planing. Phys. Med. Biol. 48, 3521–3542 (2003)
18. Tu, J.: Cross-validated Multivariate Metamodeling Methods for Physics-based Computer Simulations. In: Proceedings of the IMAC-XXI (2003)
19. Vaz, A.I.F., Vicente, L.N.: A particle swarm pattern search method for bound constrained global optimization. J. Global Optim. 39, 197–219 (2007)