Journal of Global Optimization
https://doi.org/10.1007/s10898-024-01400-5

On polling directions for randomized direct-search approaches: application to beam angle optimization in intensity-modulated proton therapy

H. Rocha[1] · J. Dias[2]

Received: 16 December 2022 / Accepted: 11 April 2024
© The Author(s) 2024

Abstract
Deterministic direct-search methods have been successfully used to address real-world challenging optimization problems, including the beam angle optimization (BAO) problem in radiation therapy treatment planning. BAO is a highly non-convex optimization problem typically treated as the optimization of an expensive multi-modal black-box function which results in a computationally time consuming procedure. For the recently available modalities of radiation therapy with protons (instead of photons) further efficiency in terms of computational time is required despite the success of the different strategies developed to accelerate BAO approaches. Introducing randomization into otherwise deterministic direct-search approaches has been shown to lead to excellent computational performance, particularly when considering a reduced number (as low as two) of random poll directions at each iteration. In this study several randomized direct-search strategies are tested considering different sets of polling directions. Results obtained using a prostate and a head-and-neck cancer cases confirmed the high-quality results obtained by deterministic direct-search methods. Randomized strategies using a reduced number of polling directions showed difficulties for the higher dimensional search space (head-and-neck) and, despite the excellent mean results for the prostate cancer case, outliers were observed, a result that is often ignored in the literature. While, for general global optimization problems, mean results (or obtaining the global optimum once) might be enough for assessing the performance of the randomized method, in real-world problems one should not disregard the worst-case scenario and beware of the possibility of poor results since, many times, it is only possible to run the optimization problem once. This is even more important in healthcare applications where the mean patient does not exist and the best treatment possible must be assured for every patient.

Keywords Derivative-free optimization · Direct-search · Polling · Random directions · Beam angle optimization · IMPT

✉ H. Rocha
hrocha@fe.uc.pt

J. Dias
joana@fe.uc.pt

1 University of Coimbra, CeBER, Faculty of Economics, Av. Dias da Silva, 165, 3004-512 Coimbra, Portugal

2 University of Coimbra, INESC Coimbra, Faculty of Economics, Rua Silvio Lima, 3030-290 Coimbra, Portugal

Published online: 10 May 2024

Mathematics Subject Classification 90C06 · 90C26 · 90C56 · 90C90

1 Introduction

Direct-search methods are an excellent alternative for finding the minimum of an objective function whose derivatives are not available. This work considers the use of directional direct-search, a derivative-free algorithm that uses a set of (polling) directions to evaluate the objective function on a finite number of points. At least $n + 1$ directions of a positive spanning set, i.e. a set that positively spans an n-dimensional search space, are typically used to guarantee that at least one of the polling directions is a descent direction [1]. This minimum number of function evaluations threshold at each iteration ensures frameworks with deterministic convergence but can be a burden in terms of computational time if the evaluation of the function is expensive. Recently, numerical results have shown that introducing randomization into otherwise deterministic frameworks leads to very competitive results with enhanced computational time efficiency, particularly when considering a reduced number (as low as two) of random poll directions at each iteration [13].

Deterministic direct-search methods have been successfully used to address real-world challenging optimization problems, including the beam angle optimization (BAO) problem in radiation therapy treatment planning [7, 18–21]. The objective of the BAO problem is to find the optimal number of irradiation beams and corresponding directions. In this work, the number of beams, n, is assumed to be defined *a priori* by the treatment planner as happens in clinical practice. The BAO problem can thus be interpreted as the optimal selection of n beam irradiation directions, being a highly non-convex optimization problem that is typically treated as the optimization of an expensive multi-modal black-box function which results in a computationally time consuming procedure [21]. Despite the high-quality of the results already obtained, several attempts have focused on improving computational times, including reducing the search space [7] or the use of surrogates [18]. For the recently available modalities of radiation therapy with protons (instead of photons), the number of degrees of freedom is increased, e.g. due to the availability of different levels of energy, making the already difficult task of obtaining optimal beam irradiation directions in a clinically acceptable time more challenging.

The use of random strategies to address a problem in a deterministic setup is becoming more common in optimization after its success in machine learning and artificial intelligence. In some domains the advantages of using random strategies are known but the question here is whether, for a deterministic setup of the optimization problem at hand, randomized methods are true competitors with deterministic ones. In this study, the performance of randomized direct-search approaches considering different sets (and different number) of poll directions are assessed using the BAO problem in intensity-modulated proton therapy (IMPT). Two different cancer cases, one prostate and one head-and-neck cases, corresponding to optimization search spaces with different dimensions, are used to compare the performance of the randomized direct-search approaches with the performance of its deterministic counterpart.

The remaining of the paper is organized as follows. In the next section we recall the main features of direct-search methods. IMPT treatment planning is introduced in Sect. 3. Section 4 describes the mathematical formulation of the BAO problem. Computational results are presented in Sect. 5 and conclusions are drawn in the last section.

2 Direct-search methods

Let us consider the unconstrained minimization of an objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Direct-search methods generate, at each iteration k, a finite number of (poll) points in the neighborhood of the current iterate $\mathbf{x}^k$. Poll points are calculated by adding terms of the form $\alpha_k \mathbf{d}$ to $\mathbf{x}^k$, where vector $\mathbf{d}$ is a direction selected from a finite set of directions $\mathbf{D_k}$ and scalar α_k is the step size. After the objective function f being evaluated at all (or some) poll points, the next iterate $\mathbf{x}^{k+1}$ is set to one poll point that improves (decreases) f and the step size is kept (or possibly increased). In case no poll point improves the objective function value, the next iterate $\mathbf{x}^{k+1}$ remains equal to the current iterate $\mathbf{x}^k$ and the step size is decreased. In this way, direct-search methods iteratively generate a sequence of non-increasing iterates $\{\mathbf{x}^k\}$. The set of directions $\mathbf{D_k}$ might be updated to obtain $\mathbf{D_{k+1}}$.

Algorithm 1 describes a direct-search method, which is typically organized around two steps at each iteration, a search step in addition to the poll step. The step size update is the most commonly adopted: keep the same step size when an improvement of the objective function is obtained (step 4) or halve the step size otherwise (step 3). The search step can (potentially) improve the performance of the method, allowing the use of any strategy or *a priori* knowledge of the problem at hand, as long as the number of tested points, $\mathbf{S_k}$, is finite. This search step is optional and can be left empty ($\mathbf{S_k} = \emptyset$). While the search step provides a more global character to the method with the possible insertion of heuristics, the more rigorous poll step allows deterministic or probabilistic convergence results.

Algorithm 1 Direct-search method

Initialization:

- Choose initial point $\mathbf{x}^0 \in \mathbb{R}^n$.
- Choose initial step size $\alpha_0 > 0$.

For k = 0, 1, 2, ...

1. Search step:
   Evaluate f at a finite number of points, $\mathbf{S_k}$.
   If $\exists\, \mathbf{x}^{k+1} \in \mathbf{S_k}$: $f(\mathbf{x}^{k+1}) < f(\mathbf{x}^k)$, select $\mathbf{x}^{k+1}$ and go to step 4.
   Otherwise, go to step 2.
2. Poll step:
   Choose a set of poll directions, $\mathbf{D_k}$.
   If $f(\mathbf{x}^k) \le f(\mathbf{x}),\ \forall \mathbf{x} \in \{\mathbf{x}^k + \alpha_k \mathbf{d_i} : \mathbf{d_i} \in \mathbf{D_k}\}$, $\mathbf{x}^{k+1} = \mathbf{x}^k$ and go to step 3.
   Otherwise, choose $\mathbf{x}^{k+1} = \mathbf{x}^k + \alpha_k \mathbf{d_i} : f(\mathbf{x}^{k+1}) < f(\mathbf{x}^k)$ and go to step 4.
3. $\alpha_{k+1} = \frac{1}{2} \times \alpha_k$.
4. $\alpha_{k+1} = \alpha_k$.

2.1 Direct-search based on deterministic descent

The convergence of direct-search methods is ensured by the poll step. If $\mathbf{D_k}$, in Algorithm 1, is a positive spanning set, then [9]:

$$\forall\, \mathbf{v} \in \mathbb{R}^n, \exists\, \mathbf{d} \in \mathbf{D_k}, \frac{\mathbf{d}^T \mathbf{v}}{\|\mathbf{d}\| \|\mathbf{v}\|} > 0. \tag{1}$$

In particular, for $\mathbf{v}$ equal to the negative gradient, $-\nabla f(\mathbf{x}_k)$, global convergence to stationary points holds [9]. This result is valid when $\mathbf{D_k}$ is chosen from a finite number of positive spanning sets. When using an infinite number of positive spanning sets, an uniform bound must be imposed in (1), considering the cosine measure of all the positive spanning sets [15].

Deterministic versions typically use the directions of a positive spanning set in the poll step to guarantee a deterministic descent method. A positive spanning set for $\mathbb{R}^n$ can be defined as a set of nonzero vectors of $\mathbb{R}^n$ whose positive combinations span $\mathbb{R}^n$. It can be shown that a positive spanning set for $\mathbb{R}^n$ contains at least $n + 1$ vectors [9]. A positive spanning set contains at least one positive basis. A positive basis is a positive spanning set that does not contain any subset that is also a positive spanning set. It can be shown that a positive basis for $\mathbb{R}^n$ contains at most $2n$ vectors [9]. Positive bases with $2n$ and $n + 1$ vectors are known as maximal and minimal positive basis, respectively. Maximal and minimal positive bases commonly used for $\mathbb{R}^n$ are $[e_1 \ \dots e_n \ -e_1 \ \dots -e_n]$ and $[e_1 \ \dots e_n \ -e]$, respectively, where e_i is the i^{th} column of the identity matrix in $\mathbb{R}^n$ and $e = [1 \ \dots \ 1]^\top$. Deterministic versions using the previous maximal and minimal positive bases are denominated in this study as **Det_2n** and **Det_n+1**, respectively.

The main feature of positive spanning sets, that motivates its use in optimization frameworks, is that unless the current iterate $\mathbf{x}^k$ is at a stationary point, there is always a direction $\mathbf{d}$ in a positive spanning set that is a descent direction (forms an acute angle with the negative gradient), i.e., there is an $\alpha > 0$ such that $f(\mathbf{x}^k + \alpha\mathbf{d}) < f(\mathbf{x}^k)$. The selection of the poll directions to use at each iteration is one of the key elements that shape a direct-search algorithm. If all poll directions of a positive spanning set are tested at each iteration, polling is called complete. Polling is called opportunistic if the first direction leading to an improvement of the objective function value is taken. In this case, the order of the poll directions at each iteration can play an important role in the computational performance of the method [3].

2.2 Direct-search based on probabilistic descent

The driving force to explore probabilistic versions of different methods is the success demonstrated by randomized approaches in machine learning and artificial intelligence. In the context of direct-search, the question that arises is whether the use of random directions (instead of polling directions from a positive spanning set) allows obtaining a descent direction often enough. Knowing that at least $n + 1$ evaluations of the objective function are required to guarantee deterministic convergence, allowing the use of a smaller number of polling directions at each iteration by probabilistic methods is particularly appealing for functions that are expensive to evaluate. Probabilistic direct-search methods consider polling directions randomly generated that may not fulfill the positive spanning property.

Deterministic direct-search methods were extended by assuming that the set of polling directions $\mathbf{D_k}$ includes only a descent direction with a certain probability [13]. Nevertheless, that probabilistic approach enjoys almost-sure global convergence (convergent with probability one) provided the polling directions $\mathbf{D_k}$ are uniformly distributed on the unit ball [13]. Furthermore, probabilistic approaches testing a reduced number (as low as two) of random poll directions at each iteration reported excellent numerical results [13]. The best theoretical choice (maximizing the probability of having a good direction among two directions uniformly drawn on the unit sphere) with best practical performance corresponds to the selection of a random vector and its negative [14].

It is worth to emphasize that Algorithm 1 can describe both deterministic and probabilistic direct-search approaches. The only difference is the choice of the set of polling directions, $\mathbf{D_k}$, in the poll step.

3 IMPT treatment planning

Radiation therapy (RT) is one of the main cancer treatments, being used in more than 50% of cancer patients, either with curative or palliative intent [2]. The goal of RT is to deliver a dose of ionizing radiation to the tumor capable of eradicating the tumor cells while trying to minimize the damage to surrounding healthy tissues which are also inevitably irradiated. Hence, the success of a treatment is closely related to the ability of delivering the prescribed dose to the tumor while simultaneously sparing, as much as possible, the healthy organs and tissues.

Radiation therapy with photon beams is clearly mainstream but the use of proton beams, in particular intensity-modulated proton therapy (IMPT), has steadily increased in clinical practice due to the advantageous depth-dose characteristics of protons: a maximum energy is deposited in a very precise location, called the Bragg peak, which contributes to obtain better trade-offs between tumor irradiation and sparing of the adjacent organs and tissues. IMPT is a sophisticated RT technique, first introduced in the late 1990s by Lomax [16], where proton beams are discretized into a large number of beamlets, i.e. narrow pencil beams, with a sequence of appropriate energies. The intensities of all beamlets are jointly optimized in a large-scale optimization problem called fluence map optimization (FMO).

Both proton and photon modalities share a common treatment planning workflow, typically resorting to a treatment planning system (TPS), to calculate the beam directions and intensities that best fulfill the clinical requirements for the treatment. RT treatment planning starts with the radiation oncologist delineating all the structures of interest on the patient's computed tomography (CT) scans: the planning target volume(s) (PTVs, that correspond to the tumor volume(s) expanded by a safety margin), and the surrounding organs-at-risk (OARs). Depending on patient specific characteristics (e.g., cancer type, stage) a medical prescription is then defined by the radiation oncologist, establishing the (prescribed) radiation doses that should be delivered to the PTV and also tolerance dose values for all OARs.

In clinical practice, a treatment planner steers a TPS in an interactive way aiming to generate a treatment plan that fulfills the prescribed and tolerance doses. The planner manually tests several different parameters that shape a (black-box) optimization objective function. Typical parameters include gantry and couch angles, structure (importance) weights and dose objectives/bounds. If the dose distribution obtained is not satisfactory, the planner updates one or more parameters, based on the outcome and mostly on experience, in an iterative process that can be very time consuming (can take several days of workload for complicated cases). The proposed dose distribution needs to be reviewed by the radiation oncologist and, in case of disapproval, the trial-and-error planning process restarts.

Optimization has been playing a decisive role in improving treatment planning optimization [4, 8, 10–12, 18]. Breedveld et al. [5] presented an excellent overview of optimization methods applied to RT planning. Automation in the planning of radiotherapy treatments is essentially based on optimization methods, aiming at releasing the planner for other important tasks, e.g. quality assurance, and tries to consistently ensure high-quality plans. While the optimization of intensities (FMO problem) has been satisfactorily solved by different optimization methods/models, the selection of the best irradiation directions—beam angle optimization (BAO) problem—still does not have a resolution method/model that simultaneously satisfies the criteria of quality of solutions and acceptable computational time. One of the reasons is the difficulty of this problem, a highly non-convex optimization problem with many local minima. In IMPT planning current practice, the number of beams and their directions are chosen subjectively based on prior trial-and-error experience and practical con-

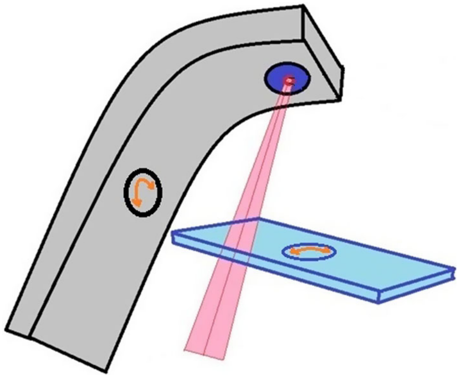

Fig. 1 Illustration of a gantry rotating around the treatment couch that can also rotate

siderations (e.g., efficiency of quality assurance and treatment delivery). However, optimal selection of beam directions can have profound impact on the quality of dose distributions and on patients' outcomes [6].

4 Beam angle optimization problem

In IMPT, proton beams exit the head of a gantry that is capable of rotating along a central axis while the patient is immobilized on a couch that can also rotate. Figure 1 illustrates the gantry rotating around the treatment couch. If the couch is fixed at zero degrees, as in Fig. 1, all beams are coplanar as they lay in the plane of rotation of the gantry. When the couch rotates, noncoplanar irradiation directions are possible. Irradiating the tumor from different beam directions allows to deliver a high dose to the tumor, towards which the different beams converge, while maintaining a lower dose outside the tumor. Optimization of beamlets, i.e. narrow pencil beams, with independent intensities allows obtaining beams with a non-uniform intensity, resulting in intensity-modulated proton therapy. The number of IMPT beams is usually 2 or 3 and rarely more than 4 or 5 [6]. Since this involves a smaller number of beams than RT treatments with photons, appropriate selection of incidence directions is even more important.

For optimization purposes, each delineated structure is discretized into voxels (small volume elements) and the radiation dose deposited in each voxel, measured in Gray (Gy), is computed using the superposition principle, i.e., considering the contribution of each beamlet. Typically, a dose matrix D, known as dose-influence matrix, mapping beamlet intensities to voxel doses is constructed, indexing the rows of D to the voxels and the columns to the beamlets. The total dose received by the voxel i, d_i, is given by $d_i = \sum_{j=1}^{N} D_{ij} w_j$, where w_j is the weight (intensity) of beamlet j. These are the decision variables when the radiation intensity optimization problem is considered. The number of voxels (V) is typically on the order of 10^6 and the number of beamlets (N) is typically on the order of 10^4 leading to large-scale optimization problems. In addition, each set of beam directions, Θ, has a different influence-dose matrix D^{Θ}. A basic radiation therapy optimization problem, for a given beam ensemble Θ is:

$$\begin{array}{lll} \min_{w^\Theta} & f(d^\Theta) & \\ s.t. & d_i^\Theta & = \sum_{j=1}^{N^\Theta} D_{ij}^\Theta w_j^\Theta, \ i = 1, \ldots, V \\ & d^\Theta & \in \ \Omega \\ & w_j^\Theta & \geq 0, \ \ j = 1, \ldots, N^\Theta, \end{array} \tag{2}$$

where w^Θ are the decision variables and dose d^Θ must be admissible, i.e., it must belong to a set Ω that can be split into several subsets corresponding to different sets of voxels (structures) that must comply with different prescription (PTVs) or tolerance (OARs) doses.

Typically, the measure used to assess the quality of a given beam ensemble, f, is the optimal value of the FMO problem [21]. In this work, the FMO objective function, f, is presented in Eq. (3):

$$f = \sum_{s=1}^{S} \sum_{i=1}^{V_s} \left\{ \underline{\lambda}_s \left[\left(L_s - \sum_{j=1}^{N_b} D_{ij} w_j \right)_+ \right]^2 + \overline{\lambda}_s \left[\left(\sum_{j=1}^{N_b} D_{ij} w_j - U_s \right)_+ \right]^2 \right\} \tag{3}$$

where S represents the set of structures considered, V_s is the number of voxels of structure s, L_s and U_s are the lower and upper bounds associated with structure s, $\underline{\lambda}_s$ and $\overline{\lambda}_s$ are the lower and upper weights for the structure s and $(\cdot)_+ = \max\{0, \cdot\}$.

In this study the optimal value of the FMO problem is calculated by matRad [22]. matRad is a research-oriented TPS, developed at the German Cancer Research Center. In this TPS, after loading a patient and selecting appropriate parameters (e.g., beam angles, objectives, constraints, etc.) it is possible to calculate the optimal solution of the FMO problem. For each beam angle set, obtaining the optimal value of the FMO problem resorting to matRad takes between one to five minutes depending on the cancer case and the number of beams considered. Thus, BAO can be simply seen as the optimization of an expensive multi-modal black-box function, f.

In contrast with RT with photons, IMPT combines a low number of proton beams incident from carefully selected directions. Furthermore, IMPT has extra degrees of freedom, provided by different levels of energy. Therefore, computational time efficiency becomes even more critical in IMPT treatment planning optimization, in particular in the selection of optimal proton beam directions (BAO) that can deeply impact the quality of dose distributions.

5 Computational results

Two different cancer cases are considered in the computational tests, corresponding to optimization problems of different dimension. The first one is a prostate cancer case where only two proton beams have to be optimized. As a noncoplanar setting was considered, the direction of each proton beam is defined by the gantry angle, θ, and the couch angle, ϕ. Thus, a four-dimensional search space is explored for the optimal selection of beam directions for the prostate cancer case. The second one is an head-and-neck cancer case where five proton beams have to be optimized, corresponding to a ten-dimensional BAO search space. The first case is less challenging (far from being easy) not only because it is smaller but also because the number of structures at stake in a prostate case is much smaller than in a head-and-neck case. The objective of this first case is, on one hand, to further validate the quality of the solutions obtained by deterministic direct-search algorithms, comparing them with a proxy of the global optimum obtained by exhaustive search. On the other hand, compare the

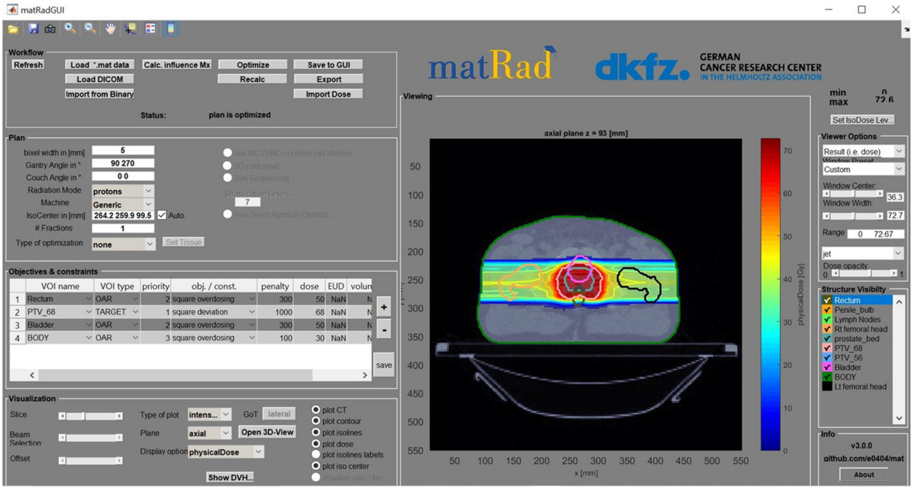

Fig. 2 IMPT for the prostate cancer case from matRad package [22]. (Color figure online)

performance of randomized strategies for different sets of poll directions with deterministic algorithms and this pseudo global optimum. In the second case, given the impossibility of using exhaustive search (it would take months of computational time), the objective is just to compare the performance of randomized strategies for different sets of poll directions with the performance of deterministic algorithms.

5.1 Prostate cancer case

The prostate cancer case included in the matRad package [22], was used in this study. The OARs included in the treatment planning optimization are the rectum and the bladder as they are near the prostate. A tolerance dose of 50 Gy is considered for these OARs. The prescribed dose for the tumor is 68 Gy. In order to avoid dose accumulation elsewhere, a structure called Body, including all the remaining normal tissue, is also included in the treatment planning optimization. The optimal FMO value for this prostate cancer case was obtained by matRad, selecting the appropriate options as displayed in Fig. 2. A common beam angle configuration for prostate IMPT corresponds to two lateral parallel opposed beams as illustrated in Fig. 2. The reasoning for this beam angle choice in the clinical setting has to do with the behavior of the traversed tissues: these are the directions that try to avoid paths with tissues that can behave differently from one session to the other due, for instance, to weight loss or inter fraction mobility. This beam angle configuration, widely used in clinical practice for prostate IMPT [6], is used as clinical benchmark.

Following the rationale that led to the choice of the clinical benchmark directions, all the (continuous) beam directions in the neighborhood of the clinical benchmark directions ($\pm 20^\circ$) are considered in the BAO procedure. Although all possible directions around the tumor can be considered, as we do in our approach, a common alternative is to consider a discretization of all possible angles with a step of five or ten degrees. Figure 3 displays the clinical benchmark directions (in red) and a discretization (using a step of five degrees) of all beam directions included in the BAO procedure (in blue). Coplanar beams (black beams) corresponds to beams with a fixed couch at zero degrees ($\phi = 0^\circ$). The total number

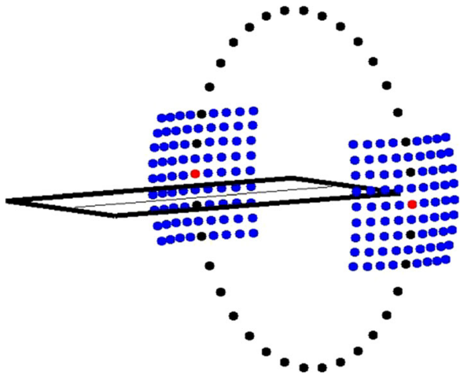

Fig. 3 Coplanar beam directions are displayed in black while a discretization of the possible noncoplanar beam directions are displayed in blue. Clinical benchmark 2-beam ensemble is displayed in red

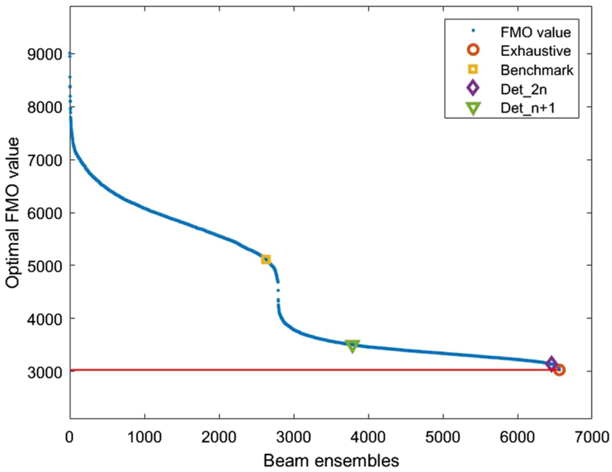

Fig. 4 Comparison of results obtained by deterministic direct-search approaches with clinical benchmark and pseudo-optimal solution

of different 2-beam ensembles is 6561. Exhaustive search testing each of these 2-beam ensembles was performed in order to obtain another benchmark that can be considered a pseudo-optimal one.

For this prostate cancer case, two deterministic approaches with opportunistic polling were tested considering a minimal and a maximal positive basis, $[e_1\ \ldots e_n\ -e]$ and $[e_1\ \ldots e_n\ -e_1\ \ldots -e_n]$, respectively. One of the objectives of this first test case was to validate the quality of BAO solutions obtained by deterministic direct-search approaches. Figure 4 presents the optimal FMO value (in descending order) for all 6561 beam ensembles calculated in the exhaustive search, as well as the solutions obtained by the deterministic direct-search using the minimal basis (**Det_n+1**) and the maximal basis (**Det_2n**). The clinical benchmark and the pseudo-optimal solution corresponding to the solution of the exhaustive search are also displayed. It can be seen that the deterministic approach **Det_2n** obtains a near-pseudo-optimal solution, validating the excellent results previously obtained with this approach [7, 18–20].

The main goal of the randomized strategies is to obtain solutions with quality similar to the solutions obtained by the deterministic strategies (in particular **Det_2n**) but in a more competitive computational time. It should be noted that although **Det_n+1** did not obtain a solution as good as **Det_2n**, it manages to obtain a solution at a level clearly better than that of the clinical benchmark which is, in itself, a very good result even if the global optimum is not reached. Deterministic approaches were compared against three randomized approaches that test a maximum (opportunistic approach) of $2n$, $n+1$, 2 and 2 symmetric (2sim) directions at each iteration, considering $\mathbf{D_k}$ constituted by: (i) **Unif** – random polling directions uniformly distributed on the unit ball [13], (ii) **Max** – polling directions randomly selected from the maximal positive basis $[e_1 \ \dots e_n \ \ -e_1 \ \dots -e_n]$, (iii) **Move2** – polling directions corresponding to the sum of two randomly selected directions of the maximal positive basis $[e_1 \ \dots e_n \ \ -e_1 \ \dots -e_n]$ (e.g., polling direction $[1 \ 1 \ 0 \ \dots \ 0]^\top$ if the first two vectors were randomly selected). Other randomized approaches were tested, including random polling directions from uniformly distributed directions by quadrants [21] or random polling directions from all linear combinations of the vectors of the maximal positive basis with coefficients in $\{0, 1\}$. Note that these last two strategies, as well as the **Unif** strategy, typically change all directions of a given beam ensemble when testing new points (beam ensembles) from the best found so far. On the other hand, the **Max** and **Move2** strategies change only one or two directions, respectively, when testing new beam ensembles from the best found so far. This strategy of only moving one (or two) directions iteratively is followed by treatment planners when performing (trial-and-error) manual planning. For the sake of better reading of the results, only the best performing strategies (**Unif**, **Max** and **Move2**) are presented. The deterministic approaches only need to be run once, since they will always lead to the same result. However, due to the random behavior of the randomized approaches, these should be run more than once since different runs can lead to different results. If the values of the objective function span over a wide interval this can be seen as a serious disadvantage for the optimization algorithm since, in real practice, most of the time the algorithm will be run only once. All randomized approaches were run twenty times.

The results of all the different approaches, including the clinical benchmark and the exhaustive search approach, are depicted in Table 1. The best minimum, mean and maximum objective function values as well as the best computational time obtained by the probabilistic approaches are highlighted with a darker cell color in the table. We can observe that all the three probabilistic approaches obtained the same minimum value, near the best theoretical value (exhaustive search) and outperforming the deterministic approaches. However, both the mean and maximum objective function values obtained by the probabilistic approaches are worse than the values obtained by the deterministic approach with the maximal basis.

Although inspection of Table 1 gives a good overview of the results, there is relevant missing information for the randomized approaches. Often only mean values are reported or, alternatively, the information of whether the optimum is obtained after a given number of runs. Although, in many cases, this is enough to assess the method performance, in real-world cases, in particular in real-world healthcare problems, it is important that each run is scrutinized. In this way, all runs of randomized approaches are reported in Figs. 5 and 6 through boxplots for the randomized approaches. Figure 5 displays the results obtained by the different approaches considering $2n$ polling directions, both in terms of optimal FMO value (5a) and in terms of computational times (5b). In terms of optimal FMO value Fig. 5a), all approaches managed to obtain excellent results and very competitive with the deterministic approach **Det_2n**. In terms of computational times (Fig. 5b) no approach outperform the others (note that clinical benchmark is almost for free as no beam angle optimization takes place). Note that the computational time of the exhaustive approach is not plotted otherwise the differences between the remaining approaches would not be visible. Despite these interesting results when using $2n$ polling directions directions, the main goal is to assess the behavior of randomized

Table 1 Computational results obtained by the different approaches for the prostate case. Best results are displayed in bold

		# of poll	f			$Time(s)$
		Directions	Min	Mean	Max	Mean
Benchmark		–	5115	5115	5115	30
Exhaustive		–	3029	3029	3029	131683
Deterministic		2n	3137	3137	3137	763
		n + 1	3503	3503	3503	310
Probabilistic		2n	**3076**	3189	3449	739
	Move2	n + 1	3089	3227	3449	542
		2	3099	3283	3531	233
		2sim	3118	3357	4845	**206**
		2n	3089	3223	3449	674
	Max	n + 1	3**076**	3216	3449	590
		2	3089	3292	3531	222
		2sim	3137	3362	4845	210
		2n	**3076**	**3168**	3449	774
	Unif	n + 1	3089	3187	**3370**	579
		2	3089	3386	5115	290
		2sim	3125	3266	3442	230

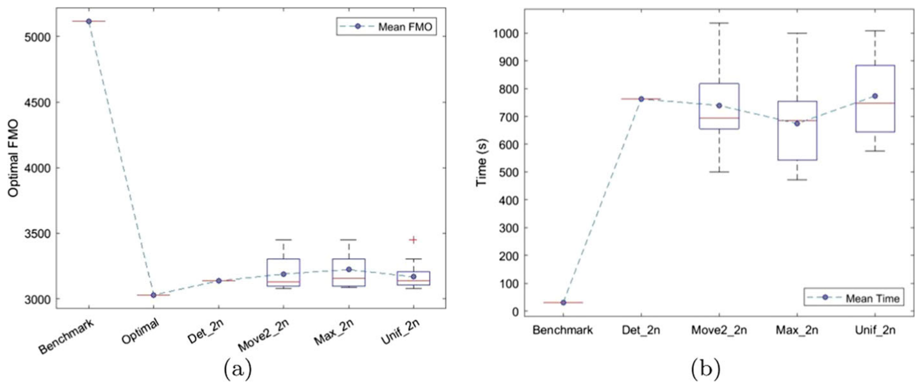

Fig. 5 Optimal FMO obtained by the different approaches considering $2n$ polling directions (**a**) and the corresponding computational times in seconds (**b**)

direct-search approaches when the number of polling directions is greatly reduced at each iteration. Figure 6 displays the results obtained by the different approaches considering only 2 polling directions. The results are in line with those reported in the literature, not deteriorating the performance in terms of optimal FMO value (Fig. 6a). It should be noted however, and this is not documented in the literature, that there are outliers whose quality of solutions is clearly inferior not even improving the clinical benchmark solution. In terms of computational times all randomized approaches with 2 polling directions (Fig. 6b) clearly outperform the approaches with $2n$ polling directions (Fig. 5b) which validates the interest in this type of

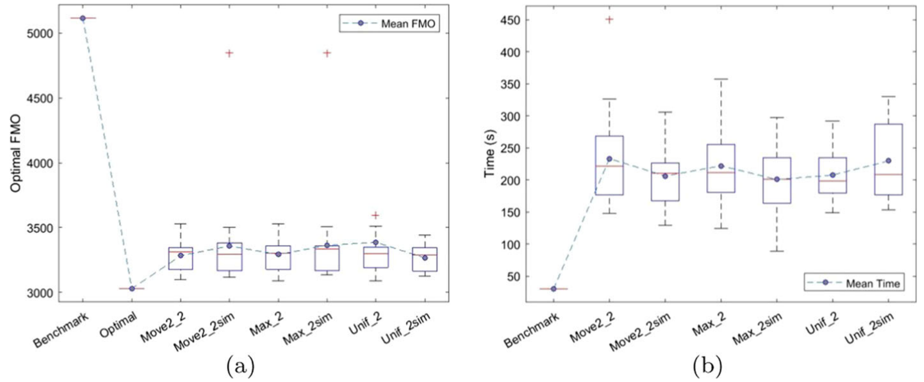

Fig. 6 Optimal FMO obtained by the different approaches considering 2 polling directions (**a**) and the corresponding computational times in seconds (**b**)

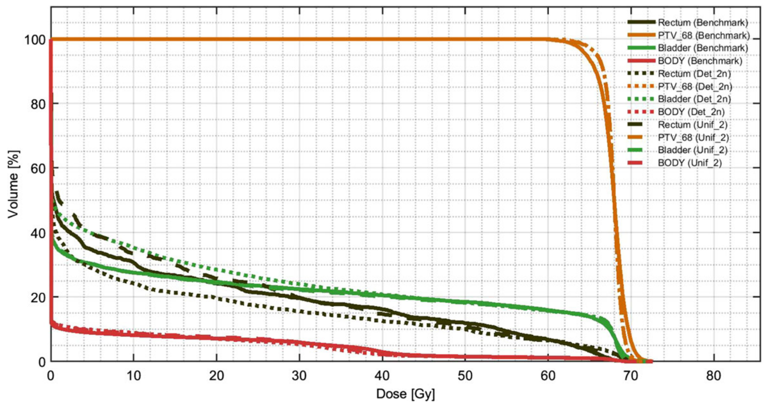

Fig. 7 Cumulative dose volume histogram comparing the results obtained by **Benchmark**, **Det_2n** and **Unif_2** treatment plans

approaches. The complete results for each of the randomized approaches (**Unif**, **Max** and **Move2**) can be found in "Appendix A".

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. DVH results for the treatment plans obtained by the benchmark solution – **Benchmark**, by the deterministic approach – **Det_2n**, and by the probabilistic approach (corresponding to a run with an average objective function value) – **Unif_2**, are displayed in Fig. 7. Both treatment plans with optimized beam directions outperform the benchmark solution. **Det_2n** have better target coverage at the expense of increased rectum sparing but decreased bladder sparing. **Unif_2** have better target coverage at the expense of rectum sparing. The dose difference occurs predominantly in the < 30 Gy portion of the DVH.

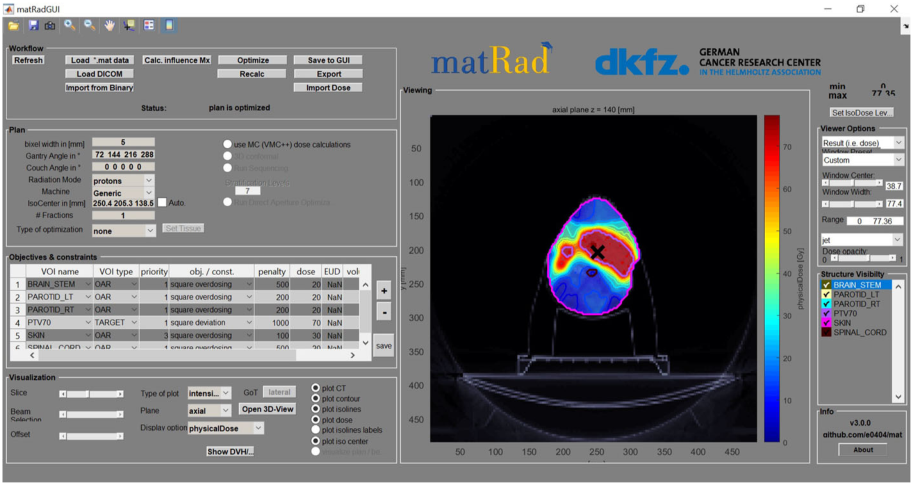

Fig. 8 IMPT for the head-and-neck cancer case from matRad package [22]. (Color figure online)

5.2 Head-and-neck cancer case

The head-and-neck cancer case included in the matRad package [22], was used in this study. Head-and-neck cancer cases are difficult cases to plan and treat because the number of sensitive OARs in that region is large. Two of the most sensitive OARs included in the treatment planning optimization are the spinal cord and the brainstem for which the tolerance dose is a maximum dose of 45 Gy and 54 Gy, respectively. Other OARs included in planning optimization are the parotid (salivary) glands which, when over-irradiated, may cause xerostomia (difficulty in swallowing) with consequences for the treatment outcome. The tolerance dose for the parotids is a mean dose of 26 Gy. The prescribed dose for the tumor is 70 Gy. The remaining normal tissue (called Skin in this case), is also included in the treatment planning optimization. The optimal FMO value for this head-and-neck cancer case was also obtained by matRad, selecting the appropriate options as displayed in Fig. 8. A common beam angle configuration for head-and-neck corresponds to coplanar equispaced beams, typically used in clinical practice. A treatment plan with five coplanar equispaced beams, as illustrated in Fig. 8, is used as clinical benchmark.

For this head-and-neck cancer case, all the (continuous) feasible (that avoid couch and gantry collision) noncoplanar beam directions are considered in the BAO procedure. Even for a sparse discretization, the total number of combinations of five beams would be so large that it would take months to run an exhaustive search. Thus, for this case, in addition to the clinical benchmark, the solutions obtained by the deterministic algorithms will be used as a benchmark given their excellent performance, not only for the prostate case but also in previous head-and-neck cases [7, 18–20]. Two deterministic approaches with opportunistic polling were tested considering a minimal and a maximal positive basis, $[e_1 \ldots e_n - e]$ and $[e_1 \ldots e_n - e_1 \ \ldots - e_n]$, respectively. Deterministic approaches were compared against the same three randomized approaches (**Unif**, **Max** and **Move2**) that test a maximum (opportunistic approach) of $2n$, $n+1$, $n/2$, 2 and 2 symmetric (2sim) directions at each iteration.

The results of all the different approaches, including the clinical benchmark, are depicted in Table 2. The best minimum, mean and maximum objective function values as well as the best computational time obtained by the probabilistic approaches are highlighted with a darker

Table 2 Computational results obtained by the different approaches for the H & N case. Best results are displayed in bold

		# of poll	f			Time (s)
		Directions	Min	Mean	Max	Mean
Benchmark		–	4965	4965	4965	60
Deterministic		2n	4125	4125	4125	5353
		n + 1	4210	4210	4210	1993
Probabilistic		2n	3937	4173	4427	4971
		n + 1	3823	4256	4748	2533
	Move2	n/2	4169	4509	4820	1205
		2	4269	4733	4965	444
		2sim	4313	4691	4918	408
		2n	3798	**4102**	**4357**	4467
		n + 1	**3780**	4349	4767	2593
	Max	n/2	4161	4576	4804	1147
		2	4483	4724	4902	461
		2sim	4258	4719	4912	423
		2n	4178	4557	4886	3587
		n + 1	4438	4663	4915	1879
	Unif	n/2	4619	4823	4965	858
		2	4636	4827	4965	341
		2sim	4661	4865	4965	**339**

cell color in the table. It is interesting to note that the best results of probabilistic approaches in terms of objective function values were always obtained considering the directions of the maximal basis. Moreover, similarly to the prostate case, the best results were obtained considering a large number of polling directions, $2n$ or $n + 1$. While the results of the probabilistic approaches are competitive with the deterministic approaches in terms of the minimum objective function value, that is not the case in terms of the mean or maximum value (or worst case), particularly when considering a reduced number of polling directions.

To complement the information given in Table 2, all the results obtained by the randomized approximations are reported in Figs. 9 and 10 through boxplots. Figure 9 displays the results obtained by the different approaches considering $2n$ polling directions, both in terms of optimal FMO value (9a) and in terms of computational times (9b). In terms of optimal FMO value (Fig. 9a), while **Max** and **Move2** approaches obtained excellent results competitive with the deterministic approach **Det_2n**, **Unif** clearly underperformed barely improving the clinical benchmark. This result is very interesting as it contradicts previous studies and shows that the choice of polling directions can be greatly influenced by the problem at hand. In this case, choosing polling directions that change one or two beams (instead of all at the same time) as done by the treatment planners seems to be a better strategy. In terms of computational times (Fig. 9b) no approach clearly outperform the others.

Although the conclusions drawn for $2n$ polling directions are interesting, considering a reduced number of polling directions at each iteration remains the focus of this study. Figure 10 displays the results obtained by the different approaches considering only 2 polling

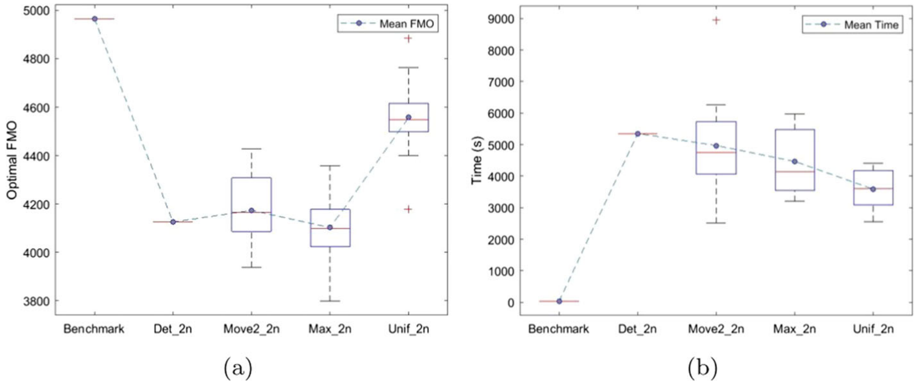

Fig. 9 Optimal FMO obtained by the different approaches considering $2n$ polling directions (**a**) and the corresponding computational times in seconds (**b**)

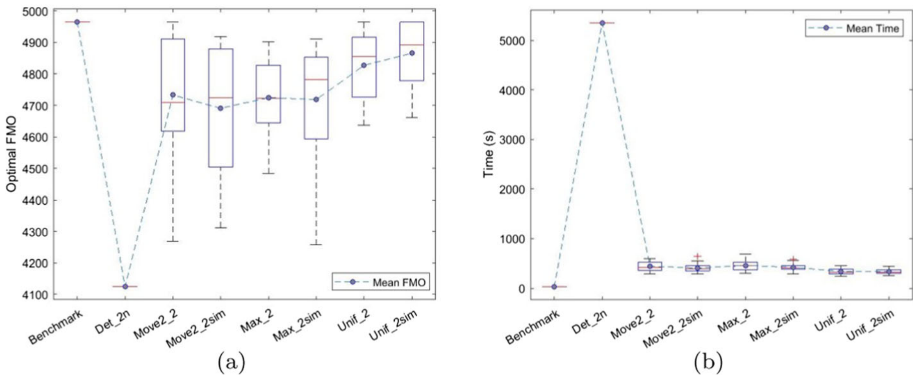

Fig. 10 Optimal FMO obtained by the different approaches considering 2 polling directions (**a**) and the corresponding computational times in seconds (**b**)

directions. In this case, it is clearly seen that the results deteriorated when a reduced number of polling directions was considered at each iteration (Fig. 10a), which, on the one hand, contradicts the results reported in the literature stating that a reduced number of polling directions obtains very competitive results and, on the other hand, it is in line with the literature that reports results deterioration for instances of increased dimensions. In terms of computational times the gains of all randomized approaches with 2 polling directions (Fig. 10b) are more pronounced also due to the increase in dimension: the difference between 2 and $2n$ increases for larger values of n. The complete results for each of the randomized approaches (**Unif**, **Max** and **Move2**) can be found in "Appendix B", being clearly visible, on one hand, the deterioration of the quality of the solutions as the number of polling directions decreases at each iteration and, on the other hand, the increasing advantage in terms of computational times for lower number of polling directions.

The existence of a marked computational time gap between the use of 2 and $2n$ polling directions gives room for testing strategies that might enhance the quality of the results obtained for the reduced number of polling directions, within the available computational budget. One of the possible strategies is to run the algorithm successively, starting each run of the algorithm at the previous optimal angular configuration. Figure 11 presents the results

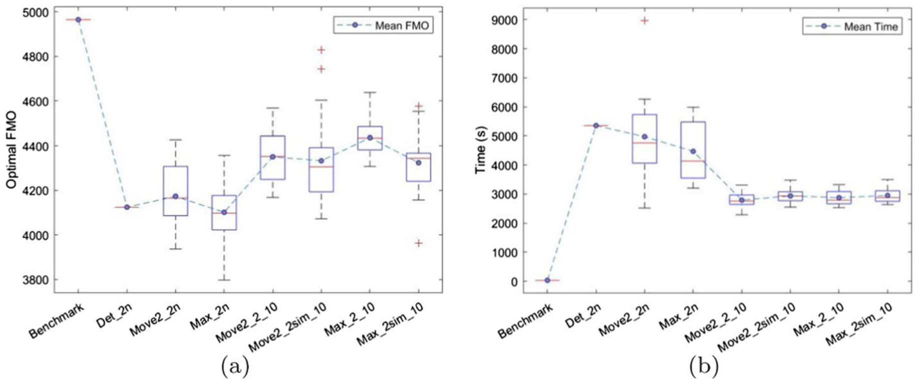

Fig. 11 Optimal FMO obtained by the different approaches considering ten runs (**a**) and the corresponding computational times in seconds (**b**)

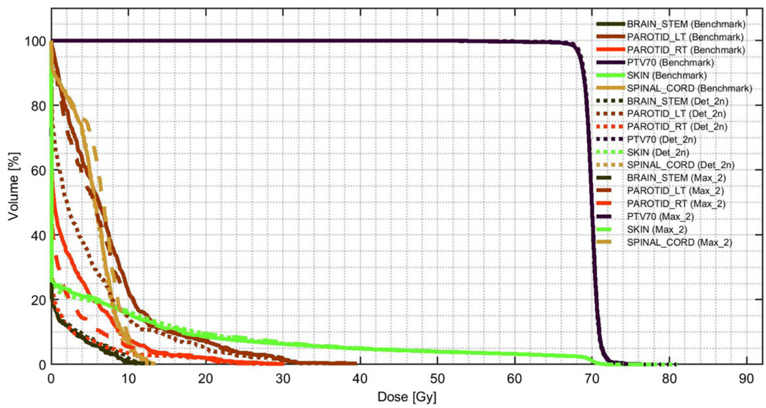

Fig. 12 Cumulative dose volume histogram comparing the results obtained by **Benchmark**, **Det_2n** and **Max_2** treatment plans

considering ten runs of **Max** and **Move2** when using only two polling directions. For benchmark purposes, the $2n$ counterparts are also displayed, as well as the deterministic approach and the clinical benchmark. Results using only two polling directions clearly improved while being still advantageous in terms of computational time. However, despite clearly outperforming now the clinical benchmark (main goal of BAO) using a reduced number of polling directions even with multiple runs falls short of the performance of the approaches that use $2n$ polling directions.

Figure 12 displays the DVH results for the treatment plans obtained by the benchmark solution—**Benchmark**, by the deterministic approach—**Det_2n**, and by the probabilistic approach (corresponding to a run with an average objective function value)—**Max_2**. Target coverage is similar for the different plans with a slight advantage for **Det_2n**. In terms of organ sparing advantage for **Det_2n** is clear.

6 Conclusions

The optimal selection of beam directions in radiation therapy is a very difficult highly non-convex optimization problem. Deterministic direct-search approaches have been successfully used to address this problem [7, 18–20]. Despite the different strategies developed to accelerate these approaches, further efficiency in terms of computational time is still required for proton radiotherapy. Given the excellent computational performance reported by direct-search approaches that use a reduced number of polling directions, in this study several randomized direct-search strategies were tested, considering different sets of polling directions. The main interest was in testing non-deterministic direct-search approaches that use a reduced number of polling directions at each iteration aiming to maximize potential gains in computational time.

Three different randomized strategies were tested, considering different sets of polling directions: **Unif**, **Max** and **Move2**. The first strategy, with good results reported in the literature [13], considers in practice the change of all angles of a given beam angular configuration, while the other two strategies correspond in practice to a change of only one or two angles at each iteration, in line with the procedure followed by the treatment planner in clinical practice. These strategies were tested using a prostate cancer case (corresponding to a 4-dimensional search space) and a head-and-neck cancer case (corresponding to a 10-dimensional search space).

The prostate cancer case confirmed the high-quality results obtained by deterministic direct-search methods, in particular when using the maximal basis [18]. The results for this case are in line with those reported in the literature, with randomized strategies obtaining very competitive results even when the number of polling directions used is reduced. However, probabilistic descent shows differences from deterministic descent as for few runs of randomized strategies the quality of solutions is clearly inferior not improving the clinical benchmark solution. While for general global optimization problems, mean results (or obtaining the global optimum once) might be enough for assessing the performance of the randomized method, in real-world problems one should look at the worst-case scenario and beware of the possibility of poor results when a reduced number of polling directions is considered at each iteration. This is even more important for healthcare problems where the mean patient does not exist and the best treatment possible must be assured for every patient and it is not possible to repeatedly run the optimization algorithm, choosing then the best solution.

The head-and-neck cancer case highlights the difficulties of direct-search approaches for larger dimensional search spaces. Even using $2n$ polling directions at each iteration, not all randomized strategies obtained high-quality results. While the randomized strategies using polling directions that change one or two angles are competitive with the deterministic approach that obtain a significant improvement with respect to the clinical benchmark, the randomized approach that uses random polling directions uniformly distributed on the unit ball behaved worst. On the other hand, reducing the number of polling directions led to a clear deterioration of the results for all the randomized approaches. As expected the computational time gap between approaches using 2 and $2n$ polling directions increased which leaves room for creative strategies that, using the computational budget of this gap, allow mitigating the loss of quality of the results.

Despite the difficulties presented by randomized strategies that use a small number of polling directions, whether due to the existence of outliers or the deterioration of results, knowledge of the problem in question can help to mitigate these setbacks. On one hand, choosing the polling directions that best fit the problem at hand can help maintain high-quality solutions even when using a reduced number of polling directions at each iteration. On the other hand, the computational gap can be used to obtain better solutions. Optimization of real-world problems, in particular healthcare problems, raises different issues, e.g. requires all solutions to be of good quality, which makes the discussion different from what is typically done in the literature and that is another reason why the conclusions are not fully aligned with those previously reported.

Despite the possible existence of outliers, the use of a reduced number of random polling directions continues to be an attractive approach to explore in future works, as long as measures are developed to mitigate this issue, e.g., risk measures. One might also consider exploring the best of deterministic and probabilistic approaches. On one hand, a numerical characteristic of deterministic directional direct-search is that most of the improvement in the objective function value occurs in the first iterations but then a long time is required to terminate with small improvements in the objective function value [17]. On the other hand, a numerical characteristic of probabilistic directional direct-search, particularly using a reduced number of polling directions, is enhanced computational times [13]. To take advantage of these characteristics, one can think of considering a hybrid method where, in the first iterations a positive spanning set is used (deterministic descent) and in the final iterations, in order to accelerate termination, a small number (two) of random polling directions is considered (probabilistic descent). Consistent quality of results and enhanced computational times can be guaranteed by the initial deterministic part and the final probabilistic part, respectively. To see if this holds in general, or if it holds only for this problem, in future work, hybrid directional direct-search approaches performance need to be assessed for a set of benchmark optimization problems.

A Prostate cancer case results obtained by the randomized approaches

See Figs. 13, 14 and 15.

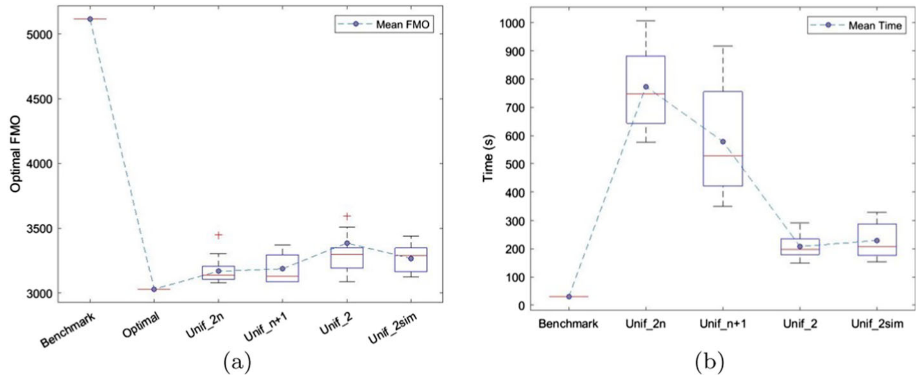

Fig. 13 Optimal FMO obtained by **Unif** (**a**) and the corresponding computational times in seconds (**b**)

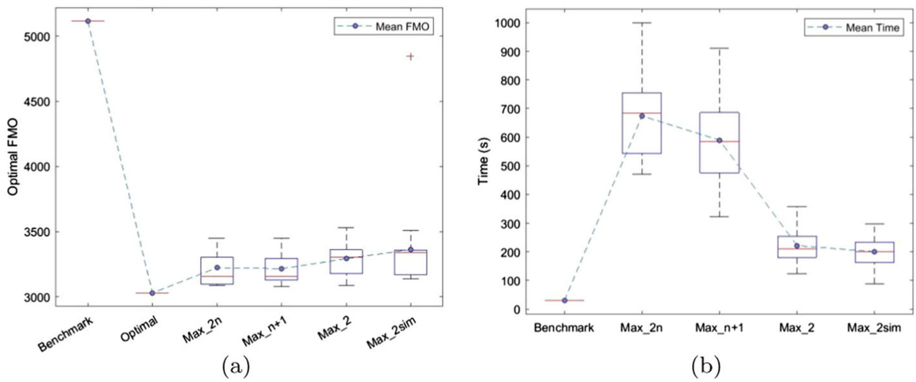

Fig. 14 Optimal FMO obtained by **Max** (**a**) and the corresponding computational times in seconds (**b**)

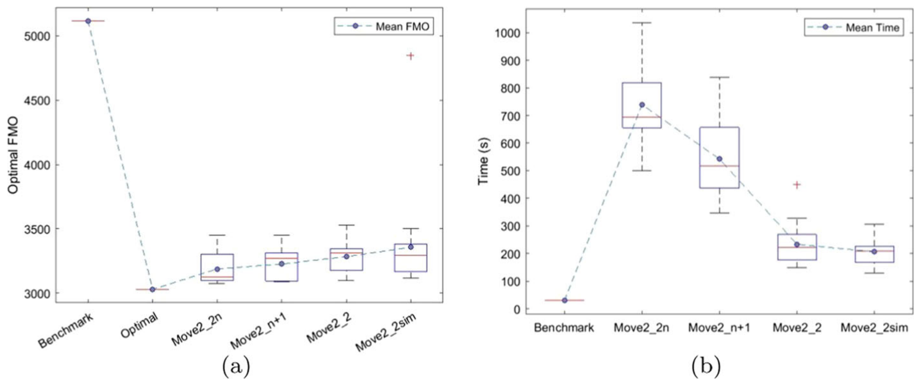

Fig. 15 Optimal FMO obtained by **Move2** (**a**) and the corresponding computational times in seconds (**b**)

B Head-and-neck cancer case results obtained by the randomized approaches

See Figs. 16, 17 and 18.

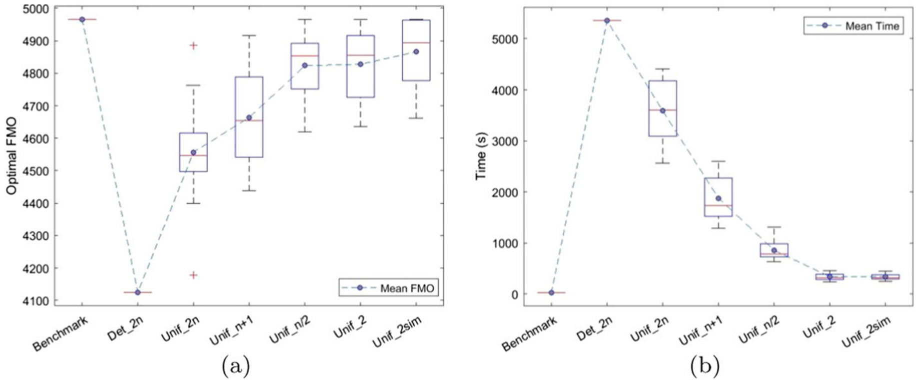

Fig. 16 Optimal FMO obtained by **Unif** (**a**) and the corresponding computational times in seconds (**b**)

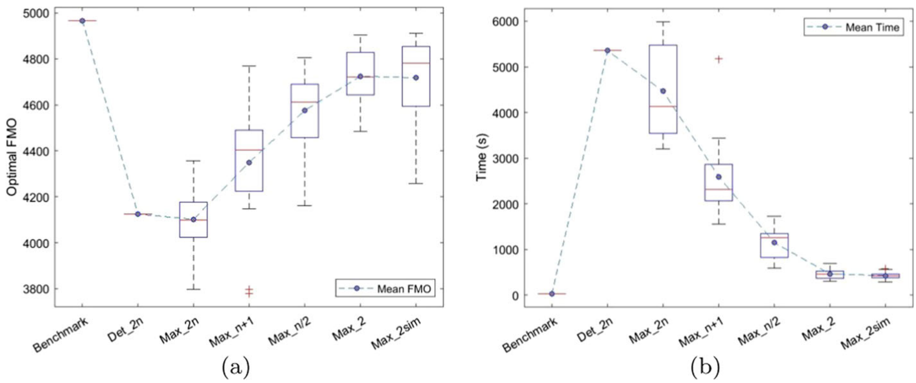

Fig. 17 Optimal FMO obtained by **Max** (**a**) and the corresponding computational times in seconds (**b**)

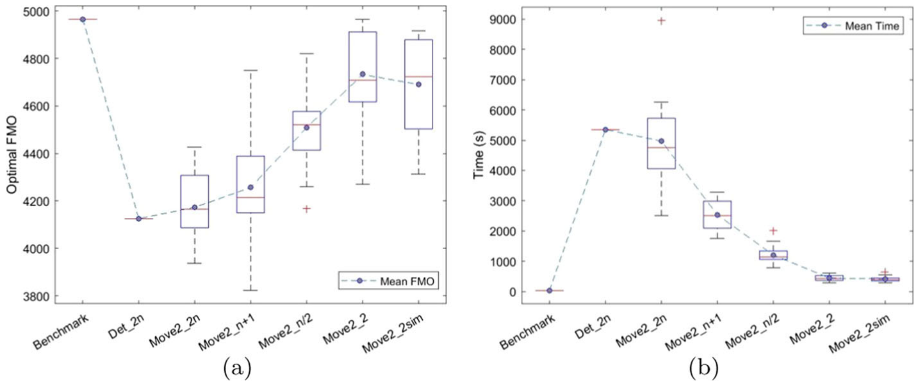

Fig. 18 Optimal FMO obtained by **Move2** (**a**) and the corresponding computational times in seconds (**b**)

Acknowledgements This work has been supported by Compete 2020 and Lisboa 2020 under project Grant POCI-01-0247-FEDER-047222 and by the Fundação para a Ciência e a Tecnologia (FCT) under project Grants UIDB/05037/2020 and UIDB/00308/2020.

Funding Open access funding provided by FCT|FCCN (b-on).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

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. American Cancer Society. Global cancer facts & figures (4th Edition). https://www.cancer.org/content/dam/cancer-org/research/cancer-factsand-statistics/global-cancer-facts-and-figures/global-cancer-facts-andfigures-4th-edition.pdf
3. Audet, C., Dennis, J.E., Jr.: Mesh adaptive direct search algorithms for constrained optimization. SIAM J. Optim. **17**, 188–217 (2006)
4. Bertsimas, D., Cacchiani, V., Craft, D., Nohadani, O.: A hybrid approach to beam angle optimization in intensity-modulated radiation therapy. Comput. Oper. Res. **40**, 2187–2197 (2013)
5. Breedveld, S., Craft, D., van Haveren, R., Heijmen, B.: Multi-criteria optimization and decision-making in radiotherapy. Eur. J. Oper. Res. **277**, 1–19 (2019)
6. Cao, W., Rocha, H., Mohan, R., Lim, G., Goudarzi, H.M., Ferreira, B.C., Dias, J.M.: Reflections on beam configuration optimization for intensity-modulated proton therapy. Phys. Med. Biol. **67**, 13TR01 (2022)
7. Carrasqueira, P., Rocha, H., Dias, J., Ventura, T., Ferreira, B.C., Lopes, M.C.: An automated treatment planning strategy for highly noncoplanar radiotherapy arc trajectories. Int. Trans. Oper. Res. **30**, 206–223 (2023)
8. Chan, T.C.Y., Craig, T., Lee, T., Sharpe, M.B.: Generalized inverse multiobjective optimization with application to cancer therapy. Oper. Res. **62**, 680–695 (2014)
9. Davis, C.: Theory of positive linear dependence. Am. J. Math. **76**, 733–746 (1954)
10. Dias, J., Rocha, H., Ferreira, B., Lopes, M.C.: A genetic algorithm with neural network fitness function evaluation for IMRT beam angle optimization. Cent. Eur. J. Oper. Res. **22**, 431–55 (2014)
11. Lim, G.J., Cao, W.: A two-phase method for selecting IMRT treatment beam angles: Branch-and-Prune and local neighborhood search. Eur. J. Oper. Res. **217**, 609–618 (2012)
12. Lim, G.J., Kardar, L., Cao, W.: A hybrid framework for optimizing beam angles in radiation therapy planning. Ann. Oper. Res. **217**, 357–383 (2014)
13. Gratton, S., Royer, C.W., Vicente, L.N., Zhang, Z.: Direct search based on probabilistic descent. SIAM J. Optim. **25**, 1515–1541 (2015)
14. Gratton, S., Royer, C.W., Vicente, L.N., Zhang, Z.: Direct search based on probabilistic feasible descent for bound and linearly constrained problems. Comput. Optim. Appl. **72**, 525–559 (2019)
15. Kolda, T.G., Lewis, R.M., Torczon, V.: Optimization by direct search: new perspectives on some classical and modern methods. SIAM Rev. **45**, 385–482 (2003)
16. Lomax, A.: Intensity modulation methods for proton radiotherapy. Phys. Med. Biol. **44**, 185–205 (1999)
17. Rocha, H., Dias, J., Ferreira, B.C., Lopes, M.C.: Pattern search methods framework for beam angle optimization in radiotherapy design. Appl. Math. Comput. **219**, 10853–10865 (2013)
18. Rocha, H., Dias, J., Ferreira, B.C., Lopes, M.C.: Selection of intensity modulated radiation therapy treatment beam directions using radial basis functions within a pattern search methods framework. J. Glob. Optim. **57**, 1065–1089 (2013)
19. Rocha, H., Dias, J., Ventura, T., Ferreira, B.C., Lopes, M.C.: A derivative-free multistart framework for an automated noncoplanar beam angle optimization in IMRT. Med. Phys. **43**, 5514–5526 (2016)
20. Rocha, H., Dias, J., Ventura, T., Ferreira, B.C., Lopes, M.C.: Beam angle optimization in IMRT: Are we really optimizing what matters? Int. Trans. Oper. Res. **26**, 908–928 (2019)
21. Rocha, H., Dias, J.: A Randomized Direct-Search Approach for Beam Angle Optimization in Intensity-Modulated Proton Therapy. In: Gervasi, O., et al. (eds.) LNCS, vol. 11621, pp. 199–214. Springer, Cham (2021)
22. Wieser, H.-P., Cisternas, E., Wahl, N., Ulrich, S., Stadler, A., Mescher, H., Muller, L.R., Klinge, T., Gabrys, H., Burigo, L., Mairani, A., Ecker, S., Ackermann, B., Ellerbrock, M., Parodi, K., Jakel, O.,

Bangert, M.: Development of the open-source dose calculation and optimization toolkit matRad. Med. Phys. **44**, 2556–2568 (2017)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.