Comparison of Hybrid Direct-Search Approaches for Beam Angle Optimization in Intensity-Modulated Proton Therapy

Humberto Rocha[1,2(✉)] and Joana Dias[1,2]

[1] University of Coimbra, FEUC, CeBER, 3004-512 Coimbra, Portugal
{hrocha,joana}@fe.uc.pt
[2] University of Coimbra, INESC-Coimbra, 3030-290 Coimbra, Portugal

Abstract. Directional direct-search methods are derivative-free optimization methods that have been used with great success in different types of real-world problems, in particular nonconvex problems with many local minima, regardless of whether derivatives are available or not. In particular, directional direct-search approaches have proven to be a good alternative to address the Beam Angle Optimization (BAO) problem – optimal selection of beam directions – in Intensity-Modulated Proton Therapy (IMPT). BAO is a highly nonconvex optimization problem that requires long computational execution times. Efficiency at this level becomes more pressing when considering protons (instead of standard photons) given the extra degrees of freedom provided by different levels of energy and the existence of different scenarios for robust IMPT plans. Despite the good performance reported in the literature when using a reduced number of random poll directions in directional direct-search methods, for the BAO problem these probabilistic (descent direction exists with a given probability) approximations showed competitive mean results against the deterministic counterpart but presented outliers, i.e., for some runs of the randomized approximation the quality of the solutions is clearly inferior. In this study we compare two hybrid strategies that combine deterministic and probabilistic directional direct-search approaches aiming at obtaining the best out of each approach: consistent quality of results and enhanced computational times, respectively. For the head-and-neck cancer case used in the computational tests, hybrid approaches obtained enhanced computational times compared to the deterministic approach with competitive results in terms of quality of the solutions obtained. In future work, hybrid directional direct-search approaches performance needs to be assessed for a set of benchmark optimization problems.

Keywords: derivative-free optimization · directional direct-search · deterministic · probabilistic · beam angle optimization · protons

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
O. Gervasi et al. (Eds.): ICCSA 2023 Workshops, LNCS 14105, pp. 467–478, 2023.
https://doi.org/10.1007/978-3-031-37108-0_30

1 Introduction

Directional direct-search methods are a class of algorithms that do not use derivatives (derivative-free) making use of a set of (polling) directions to evaluate the objective function on a finite number of points. These methods have been used with great success in different types of real-world problems, in particular nonconvex problems with many local minima, regardless of whether derivatives are available or not. Their advantage in these types of problems comes from the fact that they do not use derivatives and thus mitigate the issue of local entrapment. However, for an $n-$dimensional optimization problem, to ensure (deterministic version) that at least one polling direction is a descent direction, at least $n+1$ directions of a positive spanning set – a set of vectors (directions) that positively spans the search space – are required [1]. This minimum threshold on the number of function evaluations at each iteration guarantees frameworks with deterministic convergence at the cost of computational time efficiency particularly for functions that are expensive to evaluate. Alternatively, versions of directional direct-search algorithms that use a number of random directions smaller than n+1 (as small as two random polling directions) have been proposed in the literature, presenting good numerical results [6]. These frameworks with probabilistic convergence have computational time efficiency as their main feature.

Deterministic directional direct-search approaches have proven to be a good alternative to address the Beam Angle Optimization (BAO) problem – optimal selection of beam directions – in radiation therapy treatment planning [4,8–14]. The BAO problem is a difficult highly nonconvex optimization problem that is commonly formulated as an expensive multi-modal black-box optimization problem which results in lengthy computationally time procedures [8]. Efficiency at this level becomes more pressing when considering protons (instead of standard photons) given the extra degrees of freedom provided by different levels of energy and the existence of different scenarios for robust IMPT plans.

Despite the good performance reported in the literature when using a reduced number of random poll directions in directional direct-search methods [5,6], for the BAO problem these probabilistic (descent direction exists with a given probability) approximations showed competitive mean results against the deterministic counterpart but presented outliers, i.e., for some runs of the randomized approximation the quality of the solutions is clearly inferior [13]. Solutions for healthcare problems are required to be of high-quality every time, which explains why the conclusions in Rocha and Dias [13] are not fully aligned with those previously reported.

In this study we compare two hybrid strategies that combine deterministic and probabilistic directional direct-search approaches aiming at obtaining the best out of each approach: consistent quality of results and enhanced computational times, respectively. For the head-and-neck cancer case used in the computational tests, hybrid approaches obtained enhanced computational times compared to the deterministic approach with competitive results in terms of the quality of the solutions obtained.

The remaining of the paper is organized as follows. Next section presents the main features of directional direct-search methods. Section 3 describes IMPT treatment planning for a head-and-neck cancer case. Computational results are presented in Sect. 4 and conclusions are drawn in the last section.

2 Directional Direct-Search Methods

Directional direct-search methods make use of polling directions to generate, at each iteration k, trial points that are neighbors of the current iterate (best solution do far) $\mathbf{x}^k$. Trial points $\mathbf{x}^{k+1}$, also called poll points, are calculated as

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \alpha_k \mathbf{d},$$

where vector $\mathbf{d}$ is a polling direction selected from a set of directions $\mathbf{D_k}$ and α_k is the step size. After the objective function $f : \mathbb{R}^n \to \mathbb{R}$ being evaluated at a finite number of poll points, the next iterate $\mathbf{x}^{k+1}$ is set to one poll point that improves (decreases in a minimization problem) f. The next iterate $\mathbf{x}^{k+1}$ will remain the same as the current iterate $\mathbf{x}^k$ when no poll point improves f and in this case the step size is decreased. This iterative procedure generates a sequence of non-increasing iterates $\{\mathbf{x}^k\}$. Set $\mathbf{D_k}$ might be updated to obtain $\mathbf{D_{k+1}}$.

Directional direct-search methods are typically organized around two steps at each iteration. In the first step, called search step, a finite number of trial points $\mathbf{S_k}$ can be tested using any (rules-free) strategy or *a priori* knowledge of the problem. This step can provide a more global character to the method but is optional and can be left empty ($\mathbf{S_k} = \emptyset$). The second step, called poll step, uses polling directions to move towards regions where iteratively the objective function value is improved. Algorithm 1 depicts a direct-search method.

2.1 Deterministic Directional Direct-Search

In Algorithm 1, when considering a set of directions $\mathbf{D_k}$ that corresponds to a positive spanning set, the strict rules of the poll step ensures the (deterministic) convergence of directional direct-search methods. 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 $\mathbb{R}^n$. A positive basis is a positive spanning set with no proper set that positively spans $\mathbb{R}^n$. It can be shown that a positive basis for $\mathbb{R}^n$ contains at least $n+1$ vectors (minimal positive basis) and cannot contain more than $2n$ (maximal positive basis) [1]. Let $\mathbf{D_k}$ denote an $n \times p$ matrix whose p ($\geq n+1$) columns form a positive spanning set. One of the main features of $\mathbf{D_k}$, for optimization purposes, is that, unless the current iterate is at a stationary point, there is always a vector $\mathbf{v} \in \mathbf{D_k}$ that is a descent direction, i.e., there is an $\alpha_k > 0$ such that $f(\mathbf{x}^k + \alpha_k \mathbf{v}) < f(\mathbf{x}^k)$. Selection of the set of poll directions, $\mathbf{D_k}$, is one of the distinguishing features of a directional direct-search algorithm. When all directions of $\mathbf{D_k}$ are explored at each iteration, polling is called complete. Polling is called opportunistic when the first poll direction leading to descent is taken. In this case, the order of the poll directions may influence the computational performance of the method [2].

Algorithm 1. Directional direct-search method

Initialization:

- Set $k = 0$.
- Choose initial point $\mathbf{x}^0 \in \mathbb{R}^n$.
- Choose initial step size $\alpha_0 > 0$.
- Choose minimum step size $\alpha_{min} > 0$.

While $\alpha_k \geq \alpha_{min}$ **do**

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 $f(\mathbf{x}^k) \leq f(\mathbf{x})$, $\forall \mathbf{x} \in \{\mathbf{x}^k + \alpha_k \mathbf{d_i} : \mathbf{d_i} \in \mathbf{D_k}\}$, $\mathbf{x}^{k+1} = \mathbf{x}^k$ 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. $\alpha_{k+1} = \alpha_k$.
5. $k = k + 1$.

2.2 Probabilistic Directional Direct-Search

Randomization proved to be a winning strategy for different types of algorithms, including for deterministic directional direct-search [2]. However, when a reduced number of directions ($< n + 1$) is considered to enhance competitiveness in terms of computational time, in particular for expensive functions to evaluate, the question is no longer how to randomly choose a direction from a (positive spanning) set where a descent direction exists but rather whether a descent direction is obtained often enough. Probabilistic directional direct-search methods consider polling directions randomly generated not necessarily fulfilling the positive spanning property. Deterministic directional 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 [6]. Nevertheless, that probabilistic approach enjoys almost-sure global convergence (convergent with probability almost one) provided the polling directions $\mathbf{D_k}$ are uniformly distributed on the unit ball [6]. Furthermore, probabilistic approaches testing a reduced number (as low as two) of random poll directions at each iteration reported excellent numerical results [5,6]. However, the results reported are typically the best value obtained after a few runs or at best mean results. When looking at the results for all the runs, there may be some that are not of high-quality which is not acceptable in a healthcare context [13].

2.3 Hybrid Directional Direct-Search

Algorithm 1 can describe both deterministic and probabilistic directional direct-search approaches. The only difference is the choice, in the poll step, of the set of polling directions, $\mathbf{D_k}$, at each iteration. While the deterministic version considers $\mathbf{D_k}$ as a positive spanning set, the probabilistic version considers $\mathbf{D_k}$ as set of random directions not necessarily fulfilling the positive spanning property. The advantage of the former is that they manage to consistently obtain good results (it is a deterministic method where in two consecutive runs the same results are obtained) while the advantage of the latter is to obtain results with best efficiency in terms of computational time, in particular when the number of polling directions considered is as small as two.

One of the numerical characteristics of (deterministic) directional direct-search methods is that they typically produce most of the improvement in the objective function value in the first iterations but then take a long time to terminate with small improvements in the objective function value. Figure 1, adapted from Rocha et al. [10], illustrates this behavior considering the history of BAO procedures for four head-and-neck cancer cases. To take advantage of these characteristics, one can think of considering a hybrid method where, in the first iterations $\mathbf{D_k}$ is a positive spanning set (deterministic descent) and in the final iterations, in order to accelerate termination, a small number (two) of random polling directions is considered (probabilistic descent). This deterministic-probabilistic hybrid version has only the guarantee of convergence with probability almost one to a local minimum [6]. One could also think of an inverse hybrid version, where in the first iterations a small number (two) of random polling directions is considered (probabilistic descent) followed by iterations where $\mathbf{D_k}$ is a positive spanning set (deterministic descent). This probabilistic-deterministic hybrid version has the guarantee of convergence to a local minimum [1].

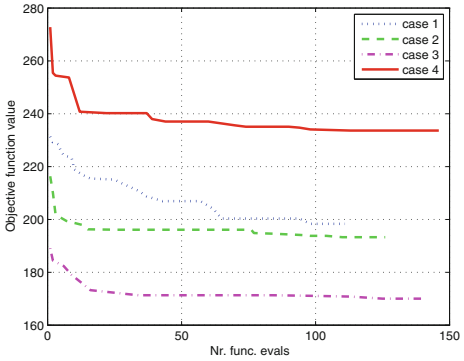

Fig. 1. Objective function value history of BAO procedures for four head-and-neck cancer cases.

3 IMPT for a Head-and-Neck Cancer Case

The different approaches are tested using a head-and-neck cancer case that is provided by matRad [16], a non-clinical radiation therapy (RT) treatment planning system written in MATLAB. Prescribed and tolerance doses for tumor volume (PTV) and organs-at-risk (OARs) considered are displayed in Table 1.

Table 1. Prescribed doses for the PTV and tolerance doses for the OARs considered.

Structure	Prescribed dose	Tolerance Dose	
		Mean	Max
PTV	70 Gy	–	–
Parotids	–	26 Gy	–
Brainstem	–	–	54 Gy
Spinal Cord	–	–	45 Gy
Skin	–	–	75 Gy

IMPT is a sophisticated RT technique, first introduced in the late 1990s by Lomax [7], 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). The optimal FMO value can be obtained by matRad, for a given set of beam directions, by selecting the appropriate options as displayed in Fig. 2. 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. 2, is used as clinical benchmark.

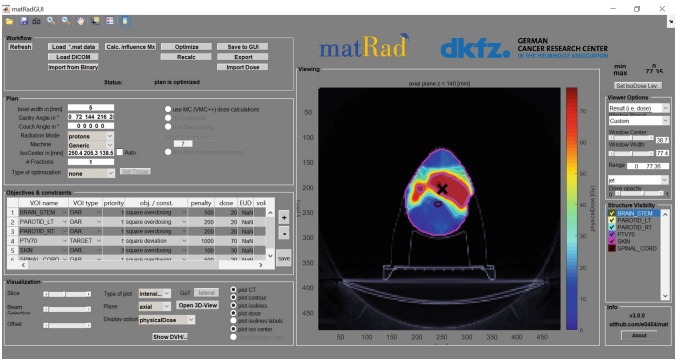

Fig. 2. IMPT for the head-and-neck case from matRad package [16].

IMPT considers a smaller number of beams than RT treatments with photons, making appropriate selection of incidence directions even more important [3]. IMPT BAO is a very difficult non-convex optimization problem, even for finding optimal sets of five beams, as is the case. Figure 3 illustrates a proton beam exiting the head of a gantry rotating around the treatment couch where the patient lies during treatment. If the couch is fixed at zero degrees, as in Fig. 3, all beams are coplanar as they lay in the plane of rotation of the gantry. When the couch rotates, noncoplanar irradiation directions are possible.

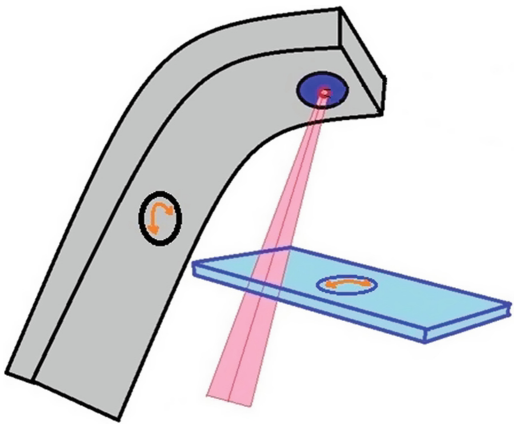

Fig. 3. Illustration of a gantry rotating around the treatment couch that can also rotate.

The IMPT BAO problem for a five-beam ensemble can be formulated as follows:

$$\min f(\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \phi_1, \phi_2, \phi_3, \phi_4, \phi_5)$$

$$\text{s.t. } (\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \phi_1, \phi_2, \phi_3, \phi_4, \phi_5) \in \mathbb{R}^{10},$$

where θ are the gantry angles, ϕ are the couch angles and f is the optimal FMO value obtained by matRad that measures the quality of the beam angle ensemble $(\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \phi_1, \phi_2, \phi_3, \phi_4, \phi_5)$. For each beam ensemble, obtaining the optimal value of the FMO problem resorting to matRad takes two to five minutes for a head-and-neck cancer case. Thus, BAO can be simply seen as the optimization of an expensive multi-modal black-box function, f.

4 Computational Results

A laptop running an Intel i7-6700 processor @ 2.60 GHz with MATLAB R2023a version was used for the computational tests. The matRad package was used to

obtain the optimal FMO value for the head-and-neck case considering five-beam ensembles. Two deterministic approaches with opportunistic polling were tested considering a minimal – $n+1$ polling directions – and a maximal – $2n$ polling directions – positive basis, $[e_1 \ \ldots e_n \ -e]$ and $[e_1 \ \ldots e_n \ -e_1 \ \ldots -e_n]$, respectively, where e_i is the ith column of the identity matrix in $\mathbb{R}^n$ and $e = [1 \ \ldots \ 1]^\top$. Deterministic directional direct-search approaches can be seen as a benchmark for the results obtained by other approaches as they have already proved to consistently obtain solutions of high-quality [15]. Probabilistic approaches were tested for a maximum (opportunistic approach) of $2n$, $n+1$, $n/2$, 2 and 2 symmetric (2sim) directions at each iteration, considering $\mathbf{D_k}$ constituted by random polling directions uniformly distributed on the unit ball [6]. Two hybrid approaches were also tested and compared. The first approach considers in the first iterations $\mathbf{D_k} = [e_1 \ \ldots e_n \ -e_1 \ \ldots -e_n]$, a maximal positive basis with $2n$ polling directions, and in the final iterations two random polling directions uniformly distributed on the unit ball. This deterministic-probabilistic hybrid version is denominated **Hyb_2n_2**. The second approach considers in the first iterations two random polling directions uniformly distributed on the unit ball, and in the final iterations $\mathbf{D_k} = [e_1 \ \ldots e_n \ -e_1 \ \ldots -e_n]$, a maximal positive basis with $2n$ polling directions. This probabilistic-deterministic hybrid version is denominated **Hyb_2_2n**. The transition in these hybrid versions is performed when the step size α_k is closer to the minimum step size α_{min} than the initial step size α_0. All approaches only perform the poll step without performing search step, i.e., $S_k = \emptyset$. Results from all approaches were compared with the five-beam clinical benchmark ensemble (denominated **Benchmark**). The deterministic approaches (denominated **Det**) only need to be run once, since they will always lead to the same result. However, due to the random behavior of the probabilistic (denominated **Unif**) and hybrid 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, are depicted in Fig. 4. Figure 4(a) shows that all different approaches, deterministic, probabilistic, and hybrid, improve the clinical benchmark solution which is automatically a very good result. However, only **Hyb_2n_2** approach manage to obtain consistently high-quality results competitive with the best deterministic approach, **Det_2n**. Moreover, computational times of **Hyb_2n_2** approach clearly improve the computational time of **Det_2n** (see Fig. 4(b)). The remaining randomized approaches, all the probabilistic ones and the probabilistic-deterministic hybrid approach, perform worst in terms of quality of solutions, with some runs that scarcely improve the clinical benchmark. However, in terms of computational time, larger gains are observed for the probabilistic approaches, in particular when using a reduced number of polling directions. Many strategies can be devised to take advantage of this computational time gap. A naïve

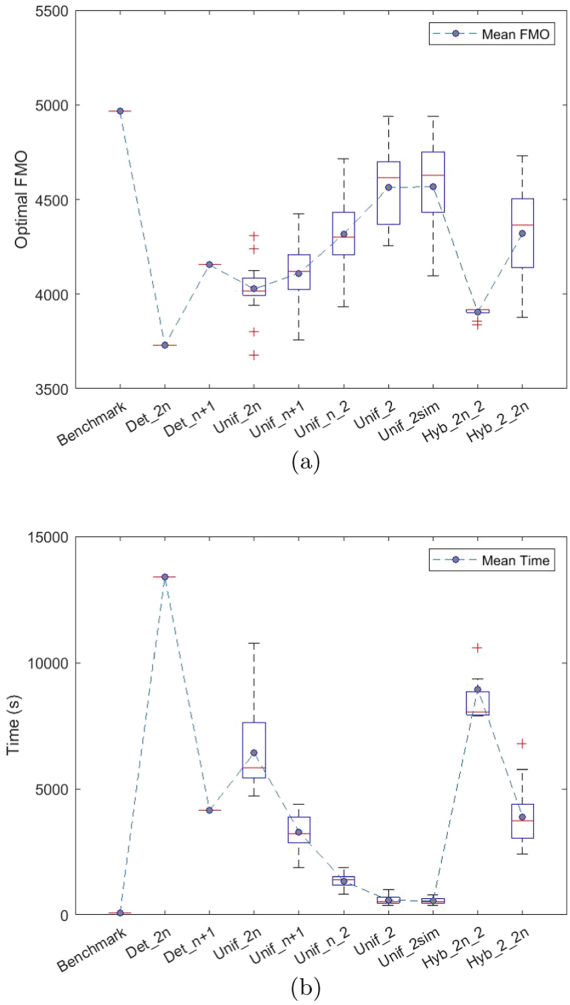

Fig. 4. Optimal FMO obtained by the different approaches – (a) and the corresponding computational times in seconds – (b).

approach would be to simply run several times the probabilistic approaches considering as starting point the best point of the last run. Figure 5 display the results when the probabilistic approaches were run ten times sequentially. While the computational time is still competitive (see Fig. 5(b)) the quality of the results improved but still fall short of the required (see Fig. 5(a)).

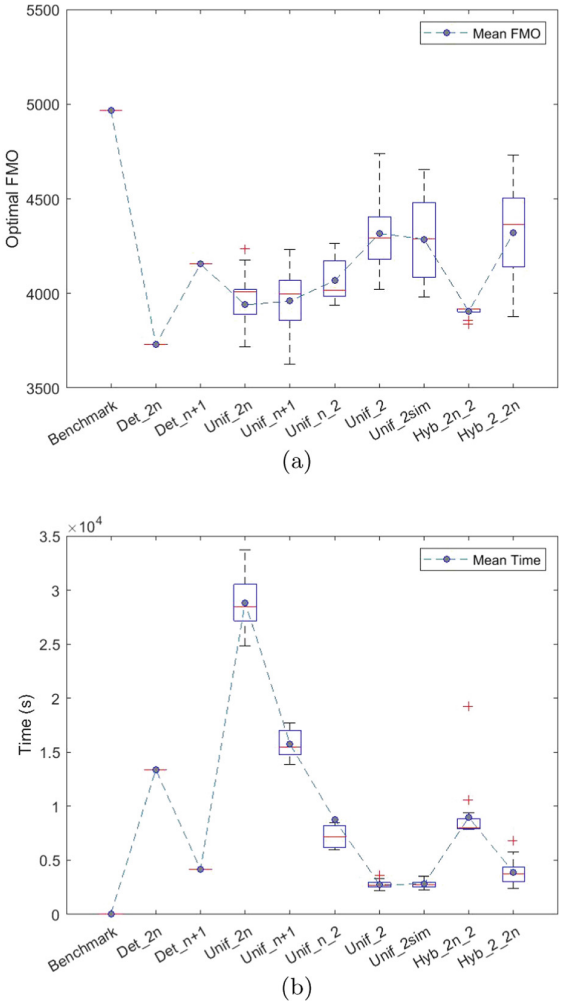

Fig. 5. Optimal FMO obtained by the different approaches – (a) and the corresponding computational times in seconds – (b).

5 Conclusions

The optimal selection of beam directions in radiation therapy is a very difficult highly nonconvex optimization problem. Deterministic directional direct-search approaches have been successfully used to address this problem. Despite the different strategies developed to accelerate these approaches, further efficiency in terms of computational time is still required for proton radiotherapy. The excellent computational performance reported by directional direct-search approaches when using a reduced number of random polling directions, showed competitive mean results for the BAO problem against the deterministic counterpart but pre-

sented outliers, i.e., for some runs of the randomized approximation the quality of the solutions is clearly inferior. In this study we compared two hybrid strategies that combine deterministic and probabilistic directional direct-search approaches aiming at obtaining the best out of each approach: consistent quality of results and enhanced computational times, respectively. The deterministic-probabilistic approach proved to be a competitive approach both in terms of consistency of quality of solutions and in terms of computational time. The main reason for the success of this approach may be related to the fact that it takes advantage of the numerical characteristics of deterministic directional direct-search methods that typically produce most of the improvement in the objective function value in the first iterations but then take a long time to terminate with small improvements in the objective function value. 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.

Acknowledgments. 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.

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. Audet, C., Dennis, J.E., Jr.: Mesh adaptive direct search algorithms for constrained optimization. SIAM J. Optim. **17**, 188–217 (2006)
3. Cao, W., et al.: Reflections on beam configuration optimization for intensity-modulated proton therapy. Phys. Med. Biol. **67**, 13TR01 (2022)
4. Carrasqueira, P., Rocha, H., Dias, J., Ventura, T., Ferreira, B., Lopes, M.C.: An automated treatment planning strategy for highly noncoplanar radiotherapy arc trajectories. Int. Trans. Oper. Res. **30**, 206–223 (2023)
5. Gratton, S., Vicente, L.N.: A merit function approach for direct search. SIAM J. Optim. **24**, 1980–1998 (2014)
6. Gratton, S., Royer, C.W., Vicente, L.N., Zhang, Z.: Direct search based on probabilistic descent. SIAM J. Optim. **25**, 1515–1541 (2015)
7. Lomax, A.: Intensity modulation methods for proton radiotherapy. Phys. Med. Biol. **44**, 185–205 (1999)
8. 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)
9. Rocha, H., Dias, J., Ferreira, B.C., Lopes, M.C.: Beam angle optimization for intensity-modulated radiation therapy using a guided pattern search method. Phys. Med. Biol. **58**, 2939–2953 (2013)
10. 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)

11. Rocha, H., Dias, J., Ferreira, B.C., Lopes, M.C.: Noncoplanar beam angle optimization in IMRT treatment planning using pattern search methods. J. Phys. Conf. Ser. **616**, 012014 (2015)
12. 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)
13. Rocha, H., Dias, J.: A randomized direct-search approach for beam angle optimization in intensity-modulated proton therapy. In: Gervasi, O., et al. (eds.) ICCSA 2021. LNCS, vol. 12953, pp. 320–332. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86976-2_22
14. 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)
15. Ventura, T., Rocha, H., Ferreira, B., Dias, J., Lopes, M.C.: Comparison of two beam angular optimization algorithms guided by automated multicriterial IMRT. Phys. Med. **64**, 210–221 (2019)
16. Wieser, H.-P., et al.: Development of the open-source dose calculation and optimization toolkit matRad. Med. Phys. **44**, 2556–2568 (2017)