A Randomized Direct-Search Approach For Beam Angle Optimization in Intensity-Modulated Proton Therapy

Humberto Rocha[1,2][0000−0002−5981−4469] and Joana Dias[1,2][0000−0003−2517−7905]

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

Abstract. Intensity-modulated proton therapy (IMPT) is a very promising alternative for radiotherapy due to the unique depth-dose characteristics of protons that allow better trade-offs between tumor irradiation and organ sparing. Optimal selection of proton beam directions – beam angle optimization (BAO) – plays a decisive role in further improving these trade-offs having a profound impact on the quality of dose distributions, particularly because in IMPT the number of beams is typically lower than in intensity-modulated radiation therapy for photons (IMRT). Computational time efficiency becomes even more critical in the optimization of proton beam directions due to the increased degrees of freedom provided by different levels of energy and the existence of different scenarios for robust IMPT plans. In this study, we consider direct-search methods to address the IMPT BAO problem given their good performance in the resolution of the IMRT BAO problem. In order to test the effectiveness of reducing the number of polling directions at each iteration, both in terms of computational time and quality of the solution, a strategy for randomly selecting a reduced number of polling directions among a set of evenly distributed directions across quadrants is proposed. This strategy considers a set of probabilistic directions, where a descent direction exists with a given probability, instead of deterministic directions that guarantee at least one descent direction. For the prostate cancer case used in the computational tests, the randomized strategy proposed shows that considering as few as two polling directions improved significantly the computational time while the resulting treatment plan is at least as good as that obtained by the deterministic method. In future work, this type of randomized approximation has to be extended and tested in different cancer cases to validate the excellent performance found for a single prostate cancer case.

Keywords: derivative-free optimization, direct-search, random directions, beam angle optimization, protons

1 Introduction

The number of cancer cases will grow by 63.1% in 2040, according to the World Health Organization [1]. More than half of all the cancer patients will need some form of radiotherapy (RT), either with curative or palliative intent. Technological advances and emergence of new treatment modalities are two of the key factors that contribute to the continuous improvement of RT treatments and make RT treatment planning an area of research constantly evolving.

The goal of RT is to eliminate the cancer cells by irradiating the tumor with a prescribed dose while sparing, as much as possible, the surrounding organs. Irradiation with photon beams is clearly mainstream in RT treatments but the use of proton beams, in particular intensity-modulated proton therapy (IMPT), presents itself as a very promising alternative due to the unique depth-dose characteristics of protons: dose is slowly deposited along the beam path before reaching a sharp peak, known as the Bragg peak, rapidly falling to almost zero beyond the peak [2]. This characteristic allows for treatment plans where a better compromise can be reached between the irradiation of the tumor and the inevitable radiation of adjacent structures, not possible with other treatment modalities. Nevertheless, obtaining high-quality treatment plans taking the most possible advantage of the unique characteristics of this treatment modality requires the optimization of different parameters including the optimal selection of proton beam directions.

In IMPT, the number of beams is typically lower than in intensity-modulated radiation therapy for photons (IMRT), being the selection of the beam directions even more critical. In addition to the smaller number of directions, the differentiating characteristics between protons and photons makes the selection of the irradiation directions in IMPT more complex. There are more degrees of freedom due to the availability of different levels of energy, and it is necessary to consider robustness due to the existence of different sources of uncertainty. Thus, obtaining optimal beam irradiation directions in a clinically acceptable time becomes even more important considering the existence of different possible scenarios required for robust plans.

The beam angle optimization (BAO) problem, i.e., the optimal selection of irradiation directions, is a very difficult problem because it is a highly non-convex optimization problem [3]. Typically, the measure used to compare the quality of different beam ensembles, and thus to guide the BAO search, is the optimal value of the fluence map optimization (FMO) problem [4], the problem of finding the optimal fluence intensities for each beam. Obtaining the optimal FMO solution for a given beam angle ensemble is time costly mainly because it requires a complete dose computation. Thus, the beam angle optimization problem can be seen as the optimization of an expensive multi-modal black-box function which results in a computationally time consuming procedure.

In previous works, direct-search methods proved to be suited for BAO in IMRT [5–11]. Although direct-search approaches require few function evaluations to converge, several attempts were made to further improve its computational time performance, including considering FMO surrogates [5] or reducing the search space [9]. Recently, different studies proposed direct-search approaches that use few (random) directions in each iteration, with numerical benefits but at the cost of guaranteed convergence to a local minimum [12,13]. Nevertheless, the almost-sure probabilistic convergence proved for these approaches translates into quality results in practice with faster computational time [13].

In this study, we propose a randomized direct-search method for BAO. Considering a prostate cancer case treated with IMPT, the proposed probabilistic approach obtained quality solutions compared to the ones obtained by the deterministic counterpart, in a faster computational time. The paper is organized as follows. In the next section we briefly describe deterministic and probabilistic direct-search methods. IMPT for a prostate cancer case is presented in Sect. 3. In the following section, the randomized direct-search proposed for BAO is described. Computational tests are presented in Sect. 5 and conclusions are made in the last section.

2 Direct-Search

Direct-search methods are a class of widely used derivative-free optimization algorithms and, as such, only use function values never resorting to any type of derivative. One of the most popular direct-search method is the Nelder-Mead method [14]. That is the algorithm underlying *fminsearch* in MATLAB [15]. The Nelder-Mead method is a simplex method that moves and manipulates the vertices of a simplex in $\mathbb{R}^n$, i.e., $n+1$ affinely independent points. In this work, we will focus on directional direct-search methods that use a set of directions, instead of simplices, to move to novel points when a decrease (considering minimization) in the objective function is obtained.

2.1 Deterministic Direct-Search

Deterministic direct-search methods consider a set of directions that correspond to a positive basis (or a positive spanning set). A positive basis for $\mathbb{R}^n$ is a set of directions (non-null vectors) that span $\mathbb{R}^n$ with nonnegative coefficients, but no proper subset does. A positive spanning set contains at least one positive basis [16]. A positive basis for $\mathbb{R}^n$ has at least $n+1$ directions (in this case is called minimal positive basis) and at most $2n$ directions (in this case is called maximal positive basis). The main motivation for using positive bases in directional direct-search methods is that at least one of its directions forms an acute angle with the negative gradient vector (unused and/or unknown) which means that this direction is a descent direction unless the current iterate is already a stationary point [17].

Direct-search methods evaluate the function in the neighborhood of the current iterate, x_k, at points of the form $\mathbf{x}^k + \alpha_k d_i$, where α_k is the step along directions d_i of a positive basis (or a positive spanning set) D_k. This procedure, called polling, aims to decrease the function value at the current iterate and is the core step of direct-search methods displayed in Algorithm 1. An optional step, called search step, can also be performed. In this step, a finite number of trial points S_k can be evaluated, not necessarily in the neighborhood of the current iterate. When the search step fails to improve the function value, or $S_k = \emptyset$, the polling around the current iterate takes place. When both search and poll steps fail to decrease the function value, the step size, α_k, is decreased – the most common choice is to halve the step size as displayed in step 3 of Algorithm 1. If one of the steps manage to find a point that improves the function value at the current iterate then α_k is increased or kept – the most common choice is to keep the same step size as displayed in step 4 of Algorithm 1.

Algorithm 1 Direct-search algorithm

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

Selection of the set of poll directions, D_k, is one of the distinguishing features of a direct-search method. Commonly used minimal and maximal positive bases are $[I \ -e]$, with I being the identity matrix of dimension n and $e = [1 \dots 1]^T$, and $[I \ -I]$, respectively. When all directions of D_k are explored at each iteration, polling is called complete and leads to the convergence of the gradient to zero for the whole sequence of iterates [18]. Polling is called opportunistic when the first poll direction leading to descent is taken, obtaining a subsequence of iterates where the gradient is driven to zero [19]. In this case, the order of the poll directions may influence the computational performance of the method [20].

2.2 Probabilistic Direct-Search

The main motivation for exploring probabilistic approaches for direct-search is the need of evaluating the function on at least $n + 1$ (minimal positive base) polling points to ensure the convergence of deterministic methods. For a large dimensional space (large n) and particularly for expensive (in terms of computational time) functions to evaluate, convergence might be too slow. Recent numerical experiments suggested that polling directions randomly generated not necessarily fulfilling the positive spanning property compare favorable to the traditional use of positive bases (or positive spanning sets), particularly if the number of directions is considerably less than $n+1$ (which can go down to two) [12]. Direct-search methods (Algorithm 1) were extended by assuming that the set of polling directions 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 D_k are uniformly distributed on the unit ball [13]. Thus, Algorithm 1 remains the same for probabilistic direct-search methods except D_k, in poll step, where random directions uniformly distributed on the unit ball are considered without restrictions on the number of directions (can be as low as one).

3 IMPT for a Prostate Case

The prostate case considered in this study is included in the matRad package [21], an open source RT treatment planning system written in MATLAB. The rectum and the bladder are in the vicinity of the prostate and for that reason are the organs-at-risk (OARs) included in the treatment planning optimization. The tolerance doses considered for this two OARs are mean doses of 50 Gy. The remaining normal tissue, called Body, is also included in the optimization just to certify that dose is not accumulating elsewhere. The prescribed dose for the planning target volume (PTV) - tumor plus a margin - is 68 Gy. Considering the appropriate options in matRad, as displayed in Figure 1, the fluence optimization for IMPT can be formulated as a quadratic nonlinear model that penalizes deviations from the prescribed/tolerated doses, implying that overdose or underdose may be clinically accepted at reduced levels, but are decreasingly acceptable for increased deviations from the prescribed/tolerated doses [22].

Two lateral parallel opposed beams are illustrated in Fig. 1 as they correspond to the most commonly used beam angle configuration for prostate proton therapy. For being widely used in clinical practice for prostate IMPT, this two-beam ensemble will be used as benchmark in our computational tests.

4 Randomized Direct-Search for BAO

IMPT BAO is a very challenging optimization problem that considers the determination of how many and which irradiation directions (angles) should be used in the treatment. For prostate cancer cases, appropriate beam selection is even more critical as proton therapy typically uses only a couple of beams. For that reason, optimal two-beam ensembles are aimed for the prostate case in study. The couch, where the patient lies during treatment, is also a degree of freedom. Figure 2 displays the benchmark two-beam ensemble (in red), possible coplanar beam directions (in black), when the couch is fixed at zero degrees, and possible noncoplanar beam directions (in blue), when couch is allowed to rotate.

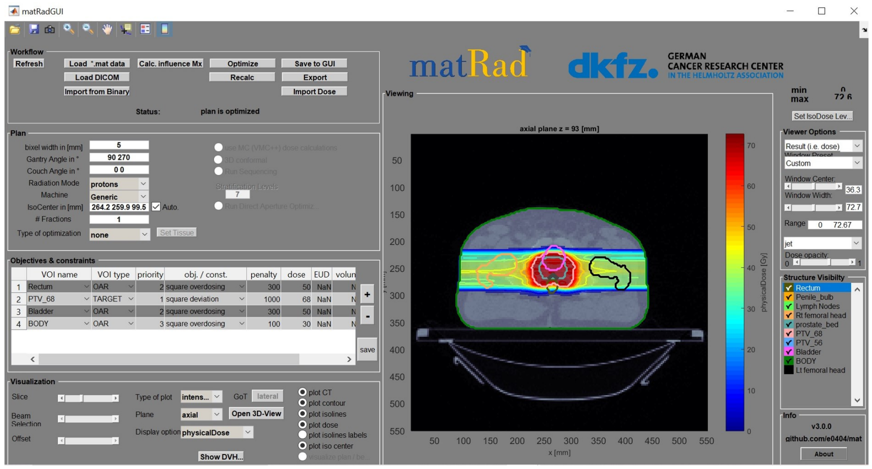

Fig. 1. IMPT for the prostate case from matRad package [21].

As highlighted before, BAO can be seen as the optimization of an expensive multi-modal black-box function, f, where the gantry angles, θ, and the couch angles, ϕ, for a two-beam ensemble are the input of f:

$$\min\ f(\theta_1, \theta_2, \phi_1, \phi_2)$$

$$\text{s.t.}\ \ (\theta_1, \theta_2, \phi_1, \phi_2) \in \mathbb{R}^4.$$

Note that this problem would be simpler if only coplanar beam directions were considered ($\phi_1 = \phi_2 = 0$) as the number of variables would be only two in a smaller search space ($\mathbb{R}^2$). The objective function $f(\theta_1, \theta_2, \phi_1, \phi_2)$ that measures the quality of the beam angle ensemble $(\theta_1, \theta_2, \phi_1, \phi_2)$ is the optimal value obtained by running the IMPT described in the previous section for each fixed set of two-beam ensembles.

We have developed deterministic direct-search approaches for BAO that were able to obtain high-quality treatment plans [5–11]. Although these approaches imply a computational time that is compatible with the clinical practice, these computational times are significant and can represent a drawback in some practical situations. The direct-search approaches for BAO we have developed consider the maximal and minimal positive bases highlighted in Sect. 2.1, which for this prostate case correspond to the directions (column-vetors) of the matrices

$$\begin{bmatrix} 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 0 & 0 & 0 & -1 \\ 0 & 1 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & 1 & -1 \end{bmatrix},$$

respectively. One of the advantages of the directions of these positive bases is that for an appropriate choice of the initial step-size (power of two) all iterates

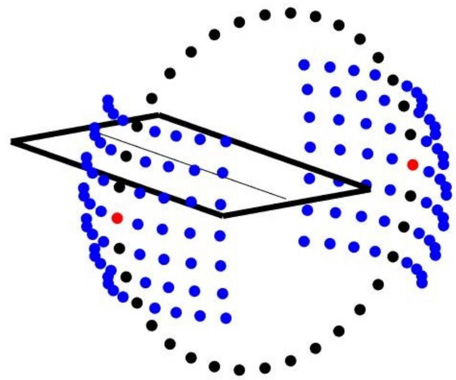

Fig. 2. Coplanar beam directions are displayed in black while some of the possible noncoplanar beam directions are displayed in blue. Benchmark 2-beam ensemble is displayed in red.

will have integer values until the step-size becomes inferior to one, which is an interesting feature for the problem at hand.

The probabilistic direct-search approach tailored for BAO has the exact same algorithm (Algorithm 1) as the deterministic one except for the set of directions D_k that drops the need to be a positive spanning set. Gratton et al. suggested polling directions D_k that are uniformly distributed on the unit ball [13]. Instead of considering the l_2-norm, that would loose the feature described in the previous paragraph, we propose the use of the l_1-norm and randomly selecting directions uniformly distributed by quadrants. Figure 3 illustrates the proposed directions for two- and three-dimensional search spaces. Note that the number of possible directions is 2^n which is equal to $2n$ for a 2-dimensional search space but will be increasingly larger than $2n$ for higher dimensional spaces.

For the prostate case considered in this study, the IMPT BAO search space is four-dimensional with possible polling directions proposed corresponding to the column-vectors of the following matrix

$$\begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 & -1 & -1 & -1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \end{bmatrix}.$$

Considering the benchmark two-beam ensemble as starting solution and an appropriate initial step-size, following any of these polling directions will always give an iterate with integer values as desired. The maximum number of random polling directions used at each iteration will, in theory, determine the pace of the algorithm. As important as verifying the computational time performance

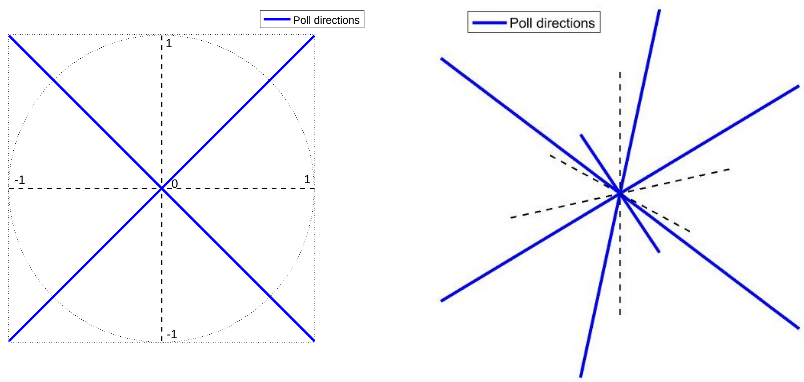

(a) (b)

Fig. 3. Poll directions, one for each quadrant, considered for two-dimensional – 3(a) and three-dimensional – 3(b) problems.

of the proposed randomized approach is to perceive the quality of the solutions obtained.

5 Computational Results

A personal computer with MATLAB R2016a version running an Intel i7-6700 processor @ 2.60 GHz was used for the computational tests. The prostate case considered is included in the matRad package, that was used for IMPT fluence optimization by selecting the appropriate options. IMPT BAO optimization was performed considering both deterministic direct-search approaches (with maximal – $2n$ polling directions – and minimal – $n+1$ polling directions) and randomized direct-search approaches (with a maximum of one, two, $n+1$, $2n$, $3n$ and $4n = 2^n$, polling directions randomly chosen at each iteration). Opportunistic polling was considered without performing search step, i.e., $S_k = \emptyset$. Results obtained were compared with the two-beam benchmark ensemble. The goal of including deterministic direct-search approaches in the computational tests is to further benchmark the results obtained by randomized approaches as results obtained by deterministic methods have already proved to be of high-quality [10]. The purpose of allowing an increased maximum number of polling directions at each iteration in randomized approaches is twofold. First, to acknowledge if more polling directions will make a difference in the quality of the solutions obtained, regardless of the computational time. Second, when considering more than $2n$ of the possible poll directions defined in Sect. 2.2, we always end up with a positive spanning set. Thus, in this case, we have a set of polling directions that is deterministically descent, i.e., there is at least one direction that is guaranteed to form an acute angle with the negative gradient vector, instead of being prob-

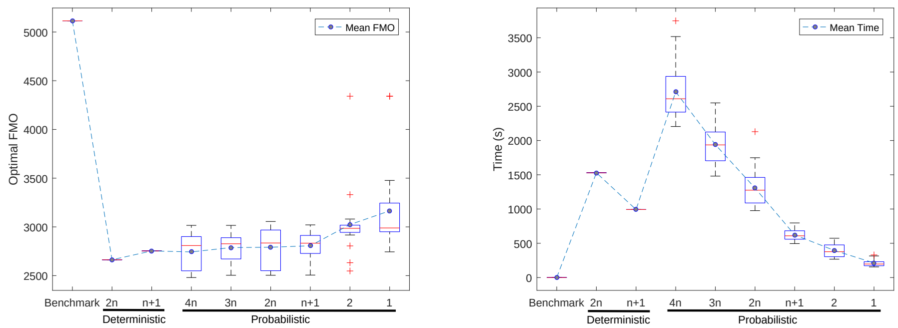

(a) (b)

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

abilistically descent which is the case when we randomly consider at most $2n$ polling directions.

As BAO is always performed resorting to the optimal FMO value, this is a natural measure to compare the quality of the solutions obtained by the different approaches. The optimal FMO value of the benchmark beam angle ensemble is 5114.82 while for the solutions of deterministic noncoplanar BAO solutions considering the maximal and the minimal positive basis are 2662.58 and 2753.35, respectively. These solutions manage to improve 48% and 46% the optimal FMO of the benchmark beam angle ensemble. The randomized approaches obtain a different solution each time the algorithm is run. For that reason, each randomized approach was run twenty times. The median (best) optimal FMO value of randomized noncoplanar BAO solutions, considering $4n$, $3n$, $2n$, $n+1$, 2 and 1 directions each iteration are 2808.52 (2480.31), 2826.79 (2503.8), 2834.77 (2503.8), 2832.83 (2505.52), 2985.26 (2547.94) and 2989.08 (2744.06), respectively. Figure 4 summarizes the performance of the different approaches both in terms of quality of solutions, as measured by the optimal FMO value, as in terms of computational times (in seconds). All BAO solutions clearly outperform the benchmark solution in terms of optimal FMO value, being deterministic BAO solutions slightly better than average randomized BAO solutions. It is interesting to see that randomized BAO solutions show no benefits from the possible inclusion of more poll directions while considering few random directions only present a small degradation of average results, that for 2 polling directions manage to obtain similar best results. In terms of computational times, the reduction of the maximum number of polling directions shows great benefits, with decreases to one third or half of the computational time when considering a maximum of 2 random directions compared to deterministic $2n$ or $n+1$ directions, respectively.

Although BAO solutions considerably improved the optimal FMO value of the benchmark solution, a set of other metrics is typically used in clinical practice to assess the quality of a treatment plan. A graphical instrument that gathers

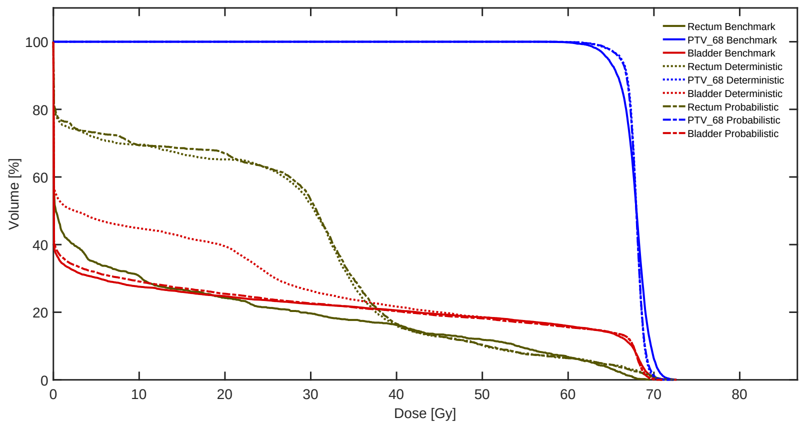

Fig. 5. Cumulative dose-volume histogram comparing the results obtained by considering the benchmark 2-beam ensemble and the 2-beam ensembles obtained by the BAO procedure with a deterministic maximal positive basis and a probabilistic selection of at most 2 poll directions at each iteration.

most of these metrics is the dose-volume histogram (DVH). The DVH displays the fraction of a structure's volume that receives at least a given dose. Ideally, the DVH line for the PTV should be at 100% volume until the prescribed dose is reached and then immediately fall to 0%, while for OARs the DVH line would ideally fall immediately to zero at 0% of the OAR volume meaning that no dose was received. Figure 5 displays the DVH results for the benchmark 2-beam ensemble, for the BAO solution obtained by the deterministic approach with $2n$ directions (the best deterministic approach in terms of optimal FMO value) and by an average BAO solution of the probabilistic approach with 2 directions (the solution with best trade-off between optimal FMO value and computational times). By simple inspection of the DVH curves, we can verify that tumor coverage is similar for both BAO solutions that clearly outperform the benchmark solution in this important feature. In terms of organ sparing, benchmark solution obtained the best results. Interestingly, while rectum sparing is similar for both BAO solutions, the deterministic approach with $2n$ directions is outperformed by the probabilistic approach with 2 directions in terms of bladder sparing.

Figure 6 displays the different two-beam ensembles whose DVHs were compared in Fig. 5. Although the polling directions followed are different as well as the maximum number of directions allowed in each iteration, it is interesting to acknowledge that the solutions obtained by deterministic and randomized methods are spatially close, which may indicate these regions as appropriate to irradiate this patient.

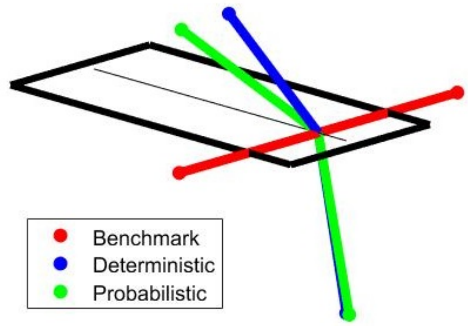

Fig. 6. Benchmark 2-beam ensemble is displayed in red while 2-beam ensembles obtained by the BAO procedure considering a deterministic maximal positive basis and a probabilistic selection of at most 2 poll directions in each iteration are displayed in blue and green, respectively.

6 Conclusions

In clinical practice, the number and directions of beams in an IMPT treatment plan are manually selected based on prior trial-and-error experience. However, the optimal selection of beam irradiation directions can deeply impact the quality of dose distributions. On one hand, the number of beams considered in IMPT is lower than in IMRT, typically 2–3 beams and rarely more than 4–5, which makes the optimal selection of beam irradiation directions more decisive. On the other hand, mainly because of high dose gradients in proton therapy, uncertainty due to anatomical variations but also other uncertainty factors specific of proton therapy, need to be addressed through robustness embedded in the optimization loop, including the optimal selection of beams. Thus, in proton therapy, decision on best beam ensembles cannot be based on dosimetric criteria alone, but must also take into consideration different sources of uncertainty. Nevertheless, assuming that robustness can be mostly handled by the FMO problem, strategies successfully developed for IMRT can be tested for IMPT.

In addition to the need for robustness, computational time becomes even more important in the optimization of irradiation directions by proton beams due to an increased number of degrees of freedom (e.g., different energy levels). In this study, we consider direct-search methods to address the IMPT BAO problem given their good performance in the resolution of the IMRT BAO problem. In order to test the effectiveness (both in terms of computational time and quality of solution) of reducing the number of polling directions at each iteration, moving from a set of deterministic directions (that guarantee at least one descent direction) to a set of probabilistic directions (where a descent direction exists with a given probability), we propose a strategy of random choice of polling directions evenly distributed across quadrants.

The proposed randomized strategy shows, for a prostate cancer case, that considering few polling directions (e.g., two) improved significantly the computational time at the cost of slightly decreasing the quality of the solution obtained. This is one of the differences from recent works on probabilistic descent that reported improved numerical behavior both in terms of computational times as well as the quality of the solution obtained when only two polling directions are considered. Nevertheless, although the optimal FMO value using only two polling directions was not the overall best, the resulting treatment plan is at least as good as that obtained by the deterministic method.

In future work, this type of randomized approximation has to be tested in more cases to validate the excellent performance found only for one prostate cancer case. Furthermore, different cancer sites where more beams are used, e.g. skull base cancer, have also to be tested to validate these approaches for optimization problems in higher dimensions. Inclusion of robustness must also be fully incorporated which was not the case in this preliminary study. Finally, different strategies for randomly selecting polling directions must be tested as well, as the success of this approach is closely linked to this choice.

Acknowledgments. This work has been supported by the Fundação para a Ciência e a Tecnologia (FCT) under project grants UTA-EXPL/FMT/0079/2019, PTDC/CCI-INF/28030/2017, UIDB/05037/2020 and UIDB/00308/2020.

References

1. Cancer tomorrow, http://gco.iarc.fr/tomorrow/home, last accessed 2021/03/17.
2. Mohan, R., Grosshans, D.: Proton therapypresent and future, Advanced drug delivery reviews **109**, 26–44 (2017)
3. Craft, D.: Local beam angle optimization with linear programming and gradient search. Phys. Med. Biol. **52**, 127135 (2007)
4. 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)
5. 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)
6. 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)
7. 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)
8. 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)
9. 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)

10. 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)
11. 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. (2021)
12. Gratton, S., Vicente, L. N.: A merit function approach for direct search. SIAM J. Optim. **24**, 1980–1998 (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. Nelder, J. A., Mead, R.: A simplex method for function minimization. Comput. J. **7**, 308–313 (1965)
15. Lagarias, J. C., Poonen, B., Wright, M. H.: Convergence of the restricted Nelder-Mead algorithm in two dimensions. SIAM J. Optim. **22**, 501–532 (2012)
16. 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)
17. Davis, C.: Theory of positive linear dependence. Am. J. Math. **76**, 733–746 (1954)
18. 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)
19. Torczon, V.: On the convergence of pattern search algorithms. SIAM J. Optim. **7**, 1–25 (1997)
20. Audet, C., Dennis Jr., J. E.: Mesh adaptive direct search algorithms for constrained optimization. SIAM J. Optim. **17**, 188–217 (2006)
21. 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)
22. 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)