

Bi-criteria Pareto optimization to balance irradiation time and dosimetric objectives in proton arc therapy

Sophie Wuyckens¹, Lewei Zhao², Michael Saint-Guillain³,
Guillaume Janssens⁴, Edmond Sterpin^{1,5}, Kevin Souris^{1,4},
Xuanfeng Ding², John A. Lee¹

¹ UCLouvain, Molecular Imaging, Radiotherapy and Oncology (MIRO), Brussels, Belgium

² Department of Radiation Oncology, Beaumont Health, Royal Oak, MI, USA

³ UCLouvain, ICTEAM, Louvain-La-Neuve, Belgium

⁴ Ion Beam Applications SA, Louvain-La-Neuve, Belgium

⁵ KULeuven, Department of Oncology, Leuven, Belgium

E-mail: sophie.wuyckens@uclouvain.be

Abstract.

Objective. Proton arc therapy (PAT) is a new delivery technique that exploits the continuous rotation of the gantry to distribute the therapeutic dose over many angular windows instead of using a few static fields, as in conventional (intensity-modulated) proton therapy. Although coming along with many potential clinical and dosimetric benefits, PAT has also raised a new optimization challenge. In addition to the dosimetric goals, the beam delivery time (BDT) needs to be considered in the objective function. Considering this bi-objective formulation, the task of finding a good compromise with appropriate weighting factors can turn out to be cumbersome. *Approach.* We have computed Pareto-optimal plans for three disease sites: a brain, a lung, and a liver, following a method of iteratively choosing weight vectors to approximate the Pareto front with few points. Mixed-Integer Programming (MIP) was selected to state the bi-criteria PAT problem and to find Pareto optimal points with a suited solver. *Main results.* The trade-offs between plan quality and beam irradiation time (*static* BDT) are investigated by inspecting three plans from the Pareto front. The latter are carefully picked to demonstrate significant differences in dose distribution and delivery time depending on their location on the frontier. The results were benchmarked against IMPT and SPArc plans showing the strength of degrees of freedom coming along with MIP optimization. *Significance.* This paper presents for the first time the application of bi-criteria optimization to the PAT problem, which eventually permits the planners to select the best treatment strategy according to the patient conditions and clinical resources available.

Keywords: Proton arc therapy, bi-objective optimization, multi-criteria optimization, Pareto front, mixed-integer programming, scalarization

Submitted to: *Phys. Med. Biol.*

1. Introduction

A common problem in radiotherapy planning is to find the best possible trade-offs between target coverage and normal tissue sparing. This is usually done by compounding a multi-objective function and manually tuning the weighting factors of each objective until the plan meets the clinical goals based on the planner’s insights and experience (Brahme 1995, Oelfke & Bortfeld 2001). This iterative trial-and-error process may turn treatment planning into a long process and sometimes it can become tedious, especially for the patients with complicated geometry with conflicting clinical goals in terms of target coverage and OARs sparing. Another popular approach is to use multi-criteria optimization (MCO), which automatically generates a set of Pareto-optimal plans for different (well-chosen) patterns of objectives weights (Küfer et al. 2000, Yu et al. 2000, Craft et al. 2005, Breedveld et al. 2019).

Pareto optimality and front. A solution (*i.e.* a treatment plan) is called *Pareto-optimal* if, by further improving one of the objectives, at least one of the other objectives gets worsened. In fact, there are as many Pareto-optimal solutions than there are possible such compromises between all the objectives at stake. By computing a set of such solutions, a so-called *Pareto front* (a.k.a. Pareto set or frontier) is obtained. In our case, the Pareto frontier represents the set of all Pareto-optimal treatment plans, over which the clinicians can efficiently navigate, and from which they can eventually pick the best trade-off depending on the clinical constraints at stake. This problem has already been extensively described in the literature for photon radiotherapy optimization such as VMAT and IMRT (Craft et al. 2012, Craft et al. 2006).

Proton arc therapy. In this article, we apply **bi-criteria** optimization to a recent innovation in proton therapy delivery, namely, proton arc therapy (PAT). PAT differs from conventional proton therapy in how the treatment is delivered to the patient, involving a sequence of many incidence angles along a rotating gantry arc, instead of using only a few fixed-angle beams. The price to pay for the larger number of irradiation angles and the increased flexibility that comes with them is that the beam delivery time (BDT) also increases, in spite of the continuous rotation of the gantry. The underlying challenge raised by this new modality is therefore to minimize the BDT concurrently with the traditional clinical or dosimetric goals. Because a true delivery time model will be certainly not convex, involving a *dynamic* delivery with mechanical parameters, we choose to rather optimize the *static* beam delivery time approximated by a simplified irradiation time, denoted BDT^{\approx} . We are investigating a bi-objective formulation where we optimize simultaneously an estimation of the irradiation time, and the dose fidelity term grouping a fixed weighted combination of objectives to treat the tumor and to penalize dose in various healthy organs.

The PAT problem. We refer to the problem of finding the best PAT treatment plan (*i.e.*, optimal dose distribution with minimal delivery time) as the *PAT problem*. A solution (*i.e.* a treatment plan) is composed of a sequence of irradiation angles. To each angle is associated a sequence of energy layers. Finally, each energy layer is composed of a sequence of spots, each spot is associated with an intensity used to modulate the dose delivered to the patient. In other words, a solution is a big sequence of spot intensities with their x-y coordinates. Informally (more formally defined in Sec. 2.1), our PAT problem consists in finding a combination of such intensity values, such that a set of dosimetric and logistic objectives are optimized.

Related work and contributions. To the best of our knowledge, this is the first time **bi-criteria** optimization is applied to proton arc therapy in order to optimize the treatment delivery time and plan quality simultaneously. Among the few published PAT optimizers currently available, **the oncology group from UCLA was the first one presenting a simultaneous approach for energy delivery sequence and plan quality optimization** (Gu et al. 2020). Based on this paper, Zhang et al. (2022) have speed up the optimization and improved both delivery efficiency and plan quality using energy matrix regularization. The other previous published paper, however, do not directly optimize the energy sequence simultaneously with the spot intensities yet. For instance, RaySearch (RaySearch Laboratories AB, Stockholm, Sweden) has recently published a method of selecting energy layers and spots, as a proof of concept coined ELSA (Early Layer and Spot Assignment), prior to spot weight optimization (Engwall et al. 2022). Additionally, Ding et al. (2016) developed the Spot-Scanning Proton Arc (SPArc) algorithm, an iterative greedy approach based on robust optimization. Bertolet & Carabe (2020) have also developed the so-called proton monoenergetic arc therapy (PMAT) technique in which, as its name suggests, **monoenergetic partial arcs are pre-selected and thus the energy sequencing optimization is a by-product of the application of the energy selection criteria. However, their research goal was different and the selection was carefully designed to enhance the dose-averaged linear energy transfer distribution within the target.** As a consequence, we present the first bi-objective optimization study applied to proton arc therapy. The proposed framework is based on Pareto **bi-criteria** optimization (BCO) theory, using an original algorithm for computing the Pareto frontier, inspired from the method introduced in Craft et al.’s (2006).

Organization. The rest of the paper is as follows. Section 2 provides a detailed description of the proton arc therapy **bi-criteria** problem and modeling as well as a sketch of the method applied to solve this problem incorporated in an innovative binary tree search technique. Section 3 demonstrates the method on three patients with different tumor sites: a brain, a lung, and a liver. Sections 4 and 5 discuss and conclude the paper. Appendices are also provided to define the mathematical model on which the PAT optimization is based (Appendix A) and to help the reader to understand the BCO Pareto approximation algorithm in more detail (Appendix B).

2. Methods

This section defines the necessary mathematical notations, states the **bi-criteria** PAT problem within a Mixed-Integer Programming (MIP) framework. A binary tree search technique is proposed to construct an approximation of the Pareto front.

2.1. Problem statement

The PAT problem pools two different types of objectives. The first objective, namely, the dose fidelity term $f(\mathbf{d})$, reflects the standard clinical goals and it is formulated with the deviation between the delivered and prescribed doses of each voxel. The second objective measures the *static* beam delivery time $\text{BDT}(\mathbf{x})$, *i.e.* the irradiation time, that is minimized to obtain efficient PAT plans. This term depends on the plan delivery sequence which includes irradiation parameters such as spot scanning speed, energy layer switching time, and burst switching time and was modeled by Zhao et al. (2022). Given the set of spot S , we consider the following bi-objective optimization problem.

$$\min_{\mathbf{x}} \quad \{f(\mathbf{d}), \text{BDT}(\mathbf{x})\} \quad (1)$$

$$\text{s.t.} \quad \mathbf{d} = \mathcal{A}\mathbf{x} \quad (2)$$

$$\mathbf{x} \in \mathbb{R}_{\geq 0}^{|S|}, \quad (3)$$

where $\mathbf{d} \in \mathbb{R}_{\geq 0}^{|V|}$ is the vectorized dose on the $|V|$ voxels ($|\cdot|$ is the cardinality operator), \mathcal{A} is the beamlet matrix, containing $|S| \times |V|$ elementary dose values, *i.e.*, the dose deposited by each spot to each voxel computed with an in-house Monte Carlo engine, **MCsquare** (Souris et al. 2016). \mathbf{x} is the column vector of spot weights, containing $|S|$ values (namely, our set of decision variables).

There is a conflicting interplay between these two main objectives. Reducing the BDT improves the patient’s comfort, mitigates motion effects over time (and therefore treatment errors), and increases the number of patients who can be treated in a day. Yet, reducing too much the BDT can also significantly degrade the dose quality.

2.2. MIP model

The fidelity term can be modeled as a linear cost function (See equation A.8 in Appendix A) with weighted linear objective functions for penalizing over-dosing and under-dosing on target, OAR doses over the limit. This term alone leads to a linear programming (LP) problem that already demonstrated its ability for dose optimization in radiotherapy as early as in the 60’s with Bahr et al. (1968) and, more recently, with Bortfeld (1999) or Romeijn et al. (2006).

The real BDT objective, naturally *dynamic*, is not convex and this leads to difficulties for any optimizer. We choose to model and minimize the irradiation time (BDT^{\approx}) using the reasonable approximation that it takes 1 s to deliver one energy layer, 5.5 s to increase the energy, 0.6 s to decrease it and 0.2 s when same energy is kept to deliver the next layer (Liu et al. 2020). In a previous work (Wuyckens

et al. 2022), we showed that this approximation of irradiation time can be modeled as a convex linear objective provided integrality constraints resulting in a mixed-integer programming (MIP) problem statement. The MIP statement has been extended to frame the bi-objective optimization problem 1, the specific mathematical expression of which is A.1-A.6 in Appendix A.

It is worth noticing that the integrality of the PAT problem is due to the BDT^{\approx} (as well as the BDT in general). In fact, it mainly depends on whether each irradiation angle and energy layer is activated or not, introducing this temporal dimension turns our PAT problem into a discrete, integer problem. In a recent study (Wuyckens et al. 2022), we proved in that the irradiation time only is responsible for the inherently complexity (i.e., the NP-hardness) of the PAT problem.

2.3. *MIP solver*

Optimization is carried out by Gurobi 9.5 (Gurobi Optimization, LLC 2022), through its Python API, for mixed-integer programming using tuned parameters to speed up the algorithm. These tuned parameters include the specification of the algorithm used to solve the continuous model, the presolving level and number of pass limits, etc. They were obtained using the Gurobi tuning tool that automates the search for parameter settings. In short, it runs the solver multiple times on the model, choosing different initial parameter settings for each call, in a search for settings that improve run time. The tuning tool parameter was run 24 hours on a model built from a phantom case with a small target. At the end, it provided a summary showing the number of parameter sets it tried and details on a few of the best parameter sets it found. On this basis, we use the parameter set with the shortest running time for the subsequent optimizations. Gurobi solves the MIP model using a linear-programming (LP) based branch-and-bound algorithm (Lawler & Wood 1966, Clausen 2003, Morrison et al. 2016). Starting from the original MIP, it removes all integrality restrictions and solves the so-called LP relaxation of the original problem through the barrier method (Marsten et al. 1990) (tuned parameter). Next, it will start the branch-and-bound procedure and will improve the bounds (relaxation) and/or incumbent solution using heuristics. Additional techniques proper to Gurobi are included to limit the size of the branch-and-bound tree that must be explored. The termination criterion is a parameter set to the time spent solving the model in this study.

2.4. *Pareto frontier approximation*

In order to solve the bi-objective optimization problem (1), it is customary in the radiotherapy field to compute a set of trade-off solutions, which are referred to as Pareto optimal points and from which the planner can conveniently pick a suitable solution. In the current article, this is carried out by blending (or scalarizing) the multiple objectives into a single compound objective, formulated as a linear combination of the initial objectives: $w_D f(\mathbf{d}) + w_T \text{BDT}^{\approx}(\mathbf{x})$. However, the adjustment of the

weighting factors, to be used as constant parameters by the optimizer, is not trivial if one wants to obtain a good approximation of the Pareto front efficiently. We decided to implement Craft et al.’s (2006) algorithm initially applied to optimizing multiple-objective intensity-modulated radiation therapy (IMRT) convex problems. In our case, the feasibility space is a polyhedron resulting from a combination of linear constraints (aside from the integer constraint). Note that the smallest polyhedron that contains all the feasible integer points is called the “convex hull” (which results from a different, optimal, set of linear constraints: the perfect problem formulation - in fact, knowing these constraints amounts to solving the problem itself). The objective function is also linear, and therefore convex. The convexity of both the feasibility space and the objective function allows us to reuse the method to approximate the Pareto front and to best reduce its front position uncertainty. Note that due to the discrete nature of our problem (and henceforth the MIP integrality constraints), the Pareto front is here a discrete set of points.

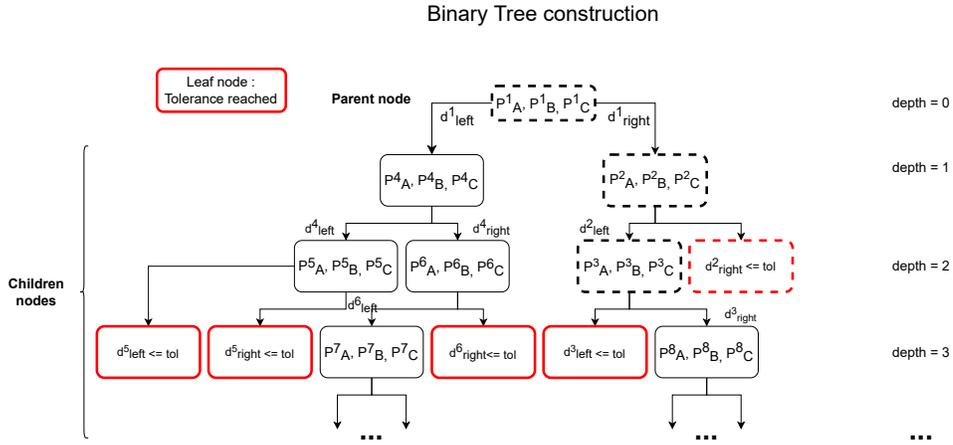


Figure 1: Pareto front construction in 2 dimensions as a binary search that produces a binary tree where each node corresponds to a point on the front. The figure contains technical details that are useful in the detailed explanation of the algorithm, in Appendix Appendix B.

For a bi-objective formulation, the BCO algorithm can be described as a search through a binary tree whose branches can be grown and explored in order to refine the Pareto front with new points wherever they are needed, such as illustrated in Figure 1. Our version of the BCO algorithm (Algorithm 1) differs from that of Craft et al.’s (2006) by the use of a binary tree, hence specialized to bi-objective optimization.

In a nutshell, the construction of the binary tree goes as follows: we start from an initial node (the *parent* node) corresponding to three points P_A , P_B and P_C designated as the *anchor points*. P_A , and P_C relate to individual minima of each objective while P_B is an intermediate solution between P_A and P_C obtained with equal weights. Next, two new points (thus corresponding to Pareto optimal solutions) can be computed, respectively, between P_A and P_B and P_B and P_C , each corresponding to *left*

and *right* branching in the binary tree depicted in Figure 1. In a 2D graph, maximal distances between segments lines $P_A - P_B$ and $P_B - P_C$ and the lower bounds of the Pareto front, *i.e.*, tangent lines through each vertex, are computed (d_{left} and d_{right} , respectively). Depending on the branching followed (the one with the minimal distance), points are renamed. For example, if one looks at left branching $P_A \rightarrow P_A$ but $P_B \rightarrow P_C$ and a new P_B is computed by solving the model using the weights average from P_A and P_C . In this way, P_B will always be the intermediate point between P_A and P_C that the optimizer needs to find. Then, new points can be subsequently and iteratively computed and added to the Pareto frontier, each time between two existing points, and so on, until some tolerance threshold (or tree depth, or time limit) is reached. However, computing each new solution (*i.e.*, a point on the Pareto front) may potentially take a significant time. In order to cope with time limitation, one must obviously prioritize the order in which the new points are computed, or in other words the order in which the branches of the binary tree are constructed, so that we end up with a representative set of points, well-balanced, to approximate the Pareto front whenever the BCO algorithm is interrupted. The algorithm, and especially the prioritization method used, is explained in details in Appendix B, including the appropriate formulas and equations.

Algorithm 1 Binary search tree

```

while  $\mathcal{D}$  and  $i \leq I_{max}$  do
  input parent ( $\mathbf{P}_A^i, \mathbf{P}_C^i$ )
  compute average weights ( $w_B^i$ ) associated with parent ( $\mathbf{P}_A^i, \mathbf{P}_C^i$ )
  optimize with weights  $w_B^i$  to find the single interior point  $\mathbf{P}_B^i$ 
  input the two presume descendants ( $\mathbf{P}_A^i, \mathbf{P}_B^i$ ) and ( $\mathbf{P}_B^i, \mathbf{P}_C^i$ )
  compute  $d_{\text{left}}^i$  and  $d_{\text{right}}^i$  and check tolerance of each
  if  $d_j^i < tol$  ( $j = \text{left}, \text{right}$ ) then
    remove  $d_j^i$  from  $\mathcal{D}$  (leaf node is reached)
  else
    update parent
  end if
  pick next node to run with largest  $d_j$  ( $j = \text{left}, \text{right}$ )
end while

```

2.5. Evaluation

As a proof of concept, we have tested the BCO algorithm on three clinical cases: a brain, a lung, and a liver. MCsquare (Souris et al. 2016) was selected as the dose calculation engine for the production of beamlets as well as the final dose computation. The sizes of each problem are reported in Table 1. The weighting factors in the dose fidelity objective were manually adjusted beforehand until the clinical goals were met in the situation where the weight of the dose fidelity term was much more important than the

BDT weight. The aim was to get rid of additional dimensions in the problem and hence to keep a bi-objective formulation for the sake of simplicity.

For each case, we report a table with the metrics of interest, i.e., the clinical goals, the irradiation time, the conformity index (CI_{RTOG} , Shaw et al. (1993)), and homogeneity index (HI, defined by the quotient of maximum and prescription dose to the target volume) for three points lying on the computed approximation of the Pareto frontier. These points are carefully picked to demonstrate the trade-off machinery such as the impact of a large weight on the dose fidelity term (big D) combined with a small weight on the BDT \approx term (small t), denoted D_t and, conversely, d_T for low dose (small d) and high time weightings (big T). An intermediate point d_t is added to the comparison. It was chosen at a glance to show a more balanced solution although the plan quality might still be questionable in some cases. Other more specific points might offer a better compromise. Only a physician could make the final decision according to each patient’s conditions. A series of plots are also provided and includes the computed Pareto fronts, the DVHs and dose distributions associated with the previously described three points, and finally the distributions of the absolute and relative weighting factors run by the BCO algorithm. In order to benchmark our results, we also evaluate an IMPT plan and a SPARC plan (Ding et al. 2016) for each disease site. They were produced in the same treatment planning system using the same solver (Gurobi) as the MIP plans, except that solely the dose fidelity term was minimized making it a linear program solved in polynomial time. Appendix C provides additional details on the dose distribution and dose-volume histograms obtained with these modalities.

	Brain	Lung	Liver
CT voxel size (mm)	$1.17 \times 1.17 \times 1.00$	$1.23 \times 1.23 \times 3.00$	$1.17 \times 1.17 \times 2.00$
Beamlet voxel size (mm)	$2.5 \times 2.5 \times 2.5$	$2.5 \times 2.5 \times 2.5$	$3 \times 3 \times 3$
nSpots	5,816	10,220	17,866
nEnergyLayers	507	762	624
nBeams	37	48	41

Table 1: Size of each optimization problem

Partial arcs were defined for each clinical case to reduce the memory load and computation time. For the brain case, one partial arc was defined. With the couch angle set to 0° , it spans angles from 40° to 130° by steps of 2.5° , resulting in 37 beams. The IMPT plan is built with three beams whose gantry angles and couch angles are respectively $95^\circ, 95^\circ, 75^\circ$ and $0^\circ, 180^\circ, 270^\circ$. The prescription to the tumor is set according to the IMPT clinical plan associated with the case, that is, 54 Gy to be delivered in 30 fractions where 95% of the tumor volume receives 100% of the dose. For this first case, beside the target coverage objectives, the sparing of the brain stem, the left optical nerve, and the optical chiasm were included in the fidelity objective given the close proximity of these OARs to the tumor.

The lung case was handled by creating two partial arcs with angles ranging from 2.5° to 197.5° and 322.5° to 357.5° , respectively, and spaced by 5° . Regarding the IMPT plan, one posterior-anterior and one lateral beams were used. The patient was prescribed a hypo-fractionnated dose of 48 Gy in 12 Gy per fraction. For this patient, optimization objectives were only set on the target coverage and irradiation time given that the dose metrics for the surrounding OARs were particularly low and uncertainties were not considered.

The third case makes use of one partial arc covering angles from 170° to 250° with a control angular step of 2° . The IMPT plan for the liver case makes use of two oblique posterior-anterior beams (couch angle = 180°) at 155° and 186° . We followed the original prescription of 67.5 Gy to be delivered in 15 fractions in total. For this patient, we have included soft objectives to control the average dose received by the liver (target volume is subtracted) and the spinal canal maximum dose level.

All the SPArc plans were initialized and splitted so that the resulting arc would reach the same sampling angle frequency and arc span (start-stop) as the MIP plans.

3. Results

We apply the proposed BCO algorithm to three clinical examples as stated in 2.5. Their results are demonstrated case by case as the following:

3.1. Brain

From Table 2 and the DVHs in Figure 2b, one can already study the balance between the objectives. The three points included in the table and shown in Figure 2a, are Pareto-optimal plans. For example, the extreme anchor point $D_{.t}$ yields a plan that has the best target coverage (D_5 - $D_{98} = 3.45$ Gy) among all plans but also the worst irradiation time. If the user wants to improve the irradiation time, it translates automatically to a degradation of the dose quality ($d_{.t}$ solution: D_5 - $D_{98} = 4.39$ Gy) and vice-versa. Although the organs at risk are always better spared with higher BDT^\approx weighting, it is only due to the fact that less beams and layers are activated in this case. For the $d_{.T}$ point, we did not pick the other extreme point (*i.e.* P_2 anchor point) because the dose distribution obtained has absolutely no clinical relevance due to too poor plan quality objective weighting versus the BDT^\approx one. The dose distributions (Figure 3) associated with each point are drastically different and emphasize the need to understand the compromises. Figure 2c shows the distribution of absolute weighting factors in decision space that were run during the BCO iterative algorithm. A color code is used to understand to which part of the Pareto front (Figure 2a) corresponds. Their distribution appears semi-continuous in the sense that the fidelity objective minimization is continuous while the BDT^\approx is discrete due to energy sequencing. Thus, this discrete behaviour impacts directly the aspect of the Pareto frontier, with few points at low dose weighting (high fidelity cost). When looking at the objective cost versus the relative

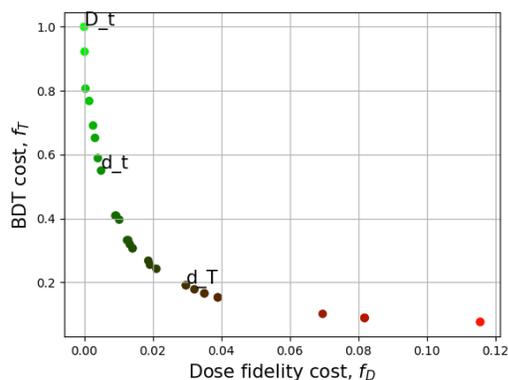
weights this time in Figure 2d, the direct point-to-point correspondence is observed with almost a symmetrical behaviour supporting the hypothesis of conflicting objectives. Regarding the benchmark (last two columns of Table 2), the MIP demonstrates its superiority compared to both IMPT and SPArc plans in terms of OARs sparing thanks to the efficient use of the additional degrees of freedom available in the optimization. The intermediate MIP plan was even capable of reducing the irradiation time compared to the IMPT plan. The SPArc plan is still faster but fails at satisfying the dose constraints.

Modality				MIP - PAT			IMPT	SPArc
				D.t	d.t	d.T		
	ROI	Metric	Limit (Gy)	Result (Gy)				
Fidelity	Target	D98	51.30	53.28	53.06	50.28	53.59	52.72
		D5	56.70	56.73	57.45	61.58	57.03	59.07
	Optical chiasm	D5	54.25	54.14	54.87	58.34	55.85	57.53
	Brain stem	Dmean	21.50	4.34	4.00	3.45	4.90	5.82
		D5	54.00	17.37	16.54	14.47	19.10	20.98
	LON	D5	54.50	54.04	55.53	54.87	54.48	56.21
	External	D5	56.65	2.31	2.12	1.10	4.18	3.27
CI				2.28	2.08	2.26	2.59	3.03
HI				1.06	1.07	1.19	1.06	1.12
Irradiation time (s)				126.2	69.4	24.1	72.2	51.40
# active spots				288	248	183	315	276
# active layers				37	26	13	40	27
# active beams				37	26	13	3	27
Optimization time (s)				18000	18000	18000	16	2155

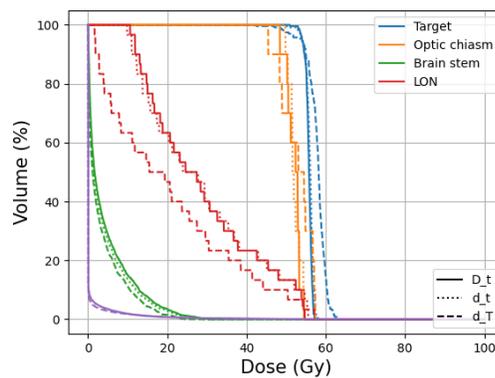
Table 2: Summary of optimization results for the brain case for three points standing on Pareto front illustrated on Figure 2a

3.2. Lung

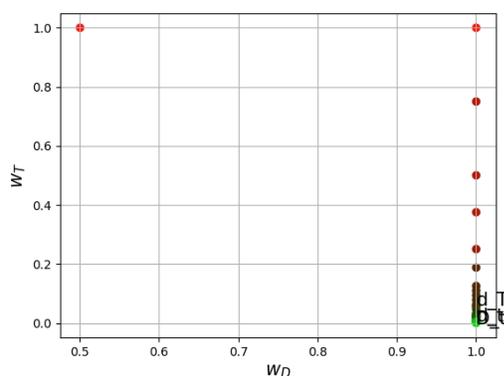
From the DVHs (Figure 4b) and metrics related (Table 3), sensitive trade-offs can be observed with the target losing rapidly its dose coverage if the time weighting gets too strong. Figure 5 illustrates the addition of beam angles to secure the dose fidelity objective and target conformity index. Compared to the brain case, this dataset is twice as large, making it harder for the MIP to find an optimal solution. This behaviour is reflected in Figure 4a where some points stack on levels, indicating the MIP was not able to push enough the optimality of the solution found in the allocated time. The IMPT and SPArc plans present significant improvement in the target coverage compared to the MIP plans. The underlying cause is that the plan parameters (spot spacing, layer spacing, and target margin) differed from those used in the MIP plans. For the lung



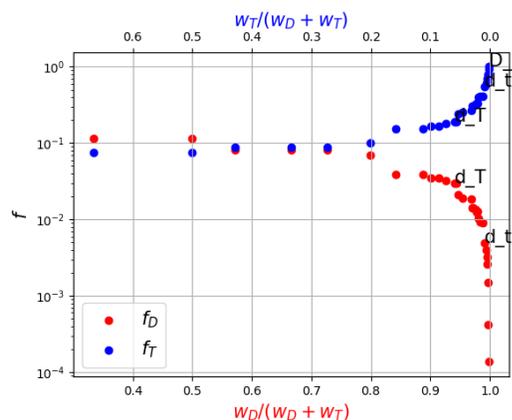
(a) Pareto front



(b) DVHs associated with 3 points on Pareto front



(c) Absolute weighting factors distribution



(d) Relative weighting factors distribution

Figure 2: BCO results obtained for the brain case. D_t is a solution where dose fidelity objective has large importance versus $BDT \approx$ objective while d_T just do the opposite and d_t is a more balanced combination of both objectives

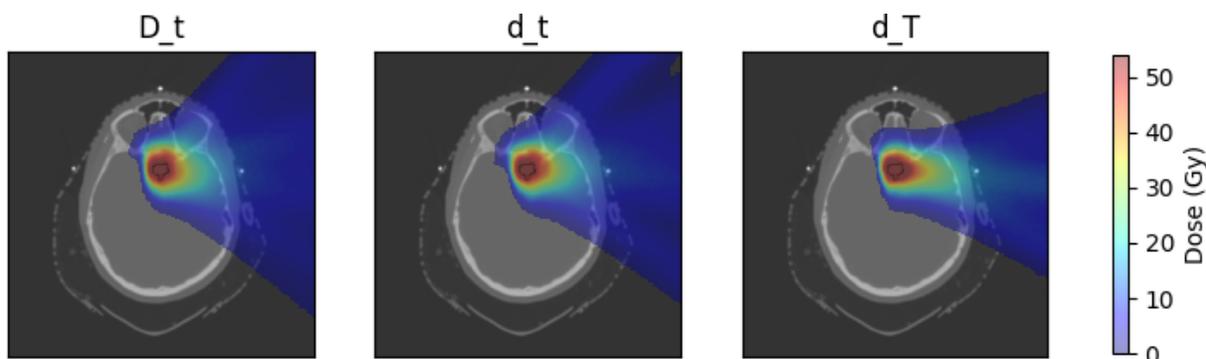


Figure 3: Brain dose distributions for three points standing on Pareto front

case, MIP plan parameters create a spot placement with insufficient coverage of the

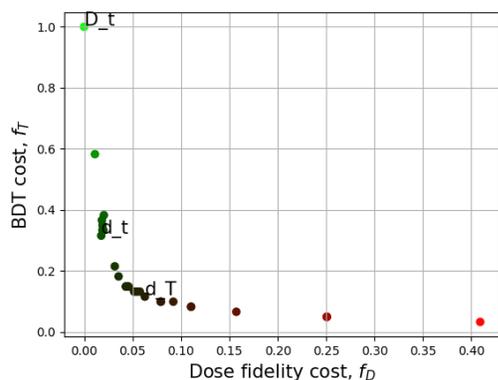
target, which the SPArc method was unable to compensate for. Hence, we increased the target margin by 1 mm, which led to better dosimetric results for SPArc. An interesting consequence is that the intermediate MIP plan (d.t) is much faster than the SPArc plan to be delivered.

Modality				MIP-PAT			IMPT	SPArc
				D.t	d.t	d.T		
	ROI	Metric	Limit (Gy)	Result (Gy)				
Fidelity	Target	D98	45.60	47.13	46.79	44.74	46.89	47.25
		D5	50.40	52.21	54.89	61.78	52.04	51.89
	Heart	D5	15.00	0.62	2.15	0.03	0.00	0.70
	Lungs-GTV	Dmean	1.00	0.68	0.68	0.83	1.07	0.99
	Chest wall	D5	30.00	2.29	2.35	1.57	1.84	3.65
	External	D5	62.00	0.06	0.05	0.03	0.02	0.06
CI				1.29	1.31	1.60	1.82	1.74
HI				1.10	1.17	1.36	1.13	1.12
Irradiation time (s)				96.6	30.5	11.2	49.7	73.80
# active spots				199	141	114	186	195
# active layers				43	17	8	29	41
# active beams				43	17	8	2	41
Optimization time (s)				18000	18000	18000	45	950

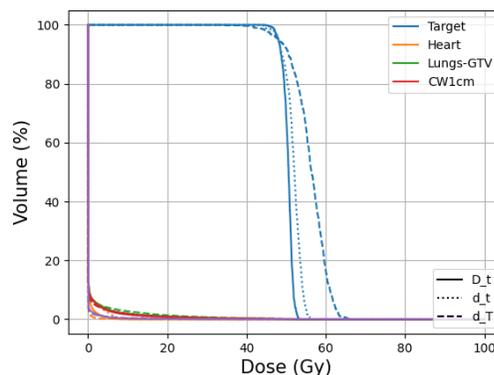
Table 3: Summary of optimization results for the lung case for three points lying on Pareto front illustrated in Figure 4a

3.3. Liver

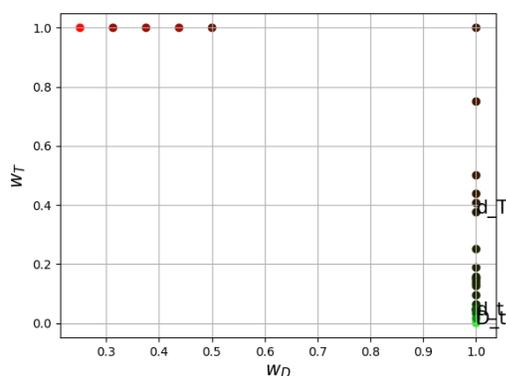
The extreme anchor point d.T point obviously leads to very poor dose quality (Figure 6b) and should not be considered at all in a real clinical study. Table 4 shows an interesting compromise where we would reduce by more than half the irradiation time and lose 2.8 Gy in the target coverage (D5-D98) when comparing points D.t and d.t. Considering the largest beamlet size, *i.e.*, lower dose resolution, the MIP optimizer was this time able to compute a very smooth approximation of the Pareto front (Figure 6a) within the time limit. Smoothness is referring in this case to the shape of the hypothetical curve connecting the solution points on the Pareto front. Figure 6d now shows clearly the inter-dependence of both objectives with a symmetric correspondence point by point. Finally, looking at the benchmark results, IMPT and SPArc plans show superior target coverage compared to the intermediate point d.T picked from the Pareto front. However, this can be explained by the fact that plan parameters had also to be tuned to obtain a sharp target DVH. The spot grid used by the MIP was too small and too sparse this time indeed. The layer spacing and spot spacing had to be both



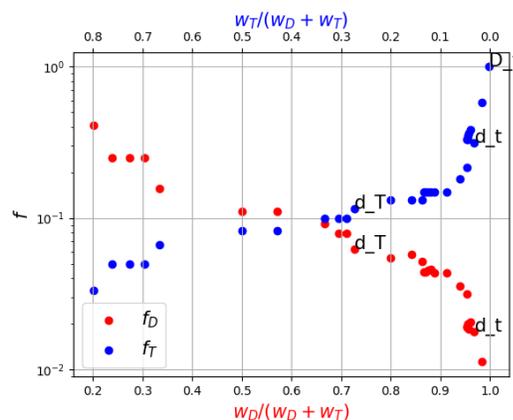
(a) Pareto front



(b) DVHs associated with 3 points on Pareto front



(c) Absolute weighting factors distribution



(d) Relative weighting factors distribution

Figure 4: BCO results obtained for the lung case. D_t is a solution where dose fidelity objective has large importance versus $BDT \approx$ objective while d_T just do the opposite and d_t is a more balanced combination of both objectives.

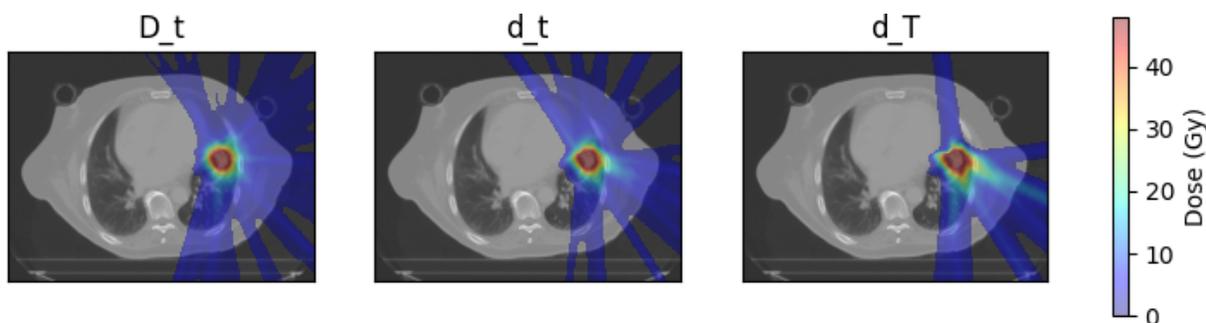


Figure 5: Lung dose distributions for three points lying on Pareto front.

reduced, whereas the target margin had to be increased, in order to bring sufficient target coverage. Appendix C discusses a few additional results to support this claim.

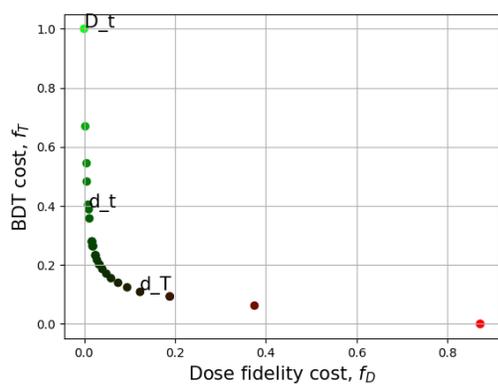
Modality				MIP - PAT			IMPT	SPArc
				D_t	d_t	d_T		
	ROI	Metric	Limit (Gy)	Result (Gy)				
Fidelity	Target	D98	64.13	66.37	66.27	60.12	66.20	66.17
		D5	70.88	71.13	73.82	99.96	72.01	71.88
	Liver-IGTV	Dmean	4.10	3.54	3.09	2.14	3.88	4.45
	Spinal Canal	D5	8.00	4.53	4.77	3.16	7.9	6.45
	External	D5	71.25	1.25	0.72	0.10	1.10	2.66
CI				1.14	1.00	0.57	1.52	1.59
HI				1.04	1.05	1.25	1.08	1.09
Irradiation time (s)				103.20	40.10	11.20	54.50	65.80
# active spots				640	557	354	680	714
# active layers				41	23	8	32	36
# active beams				41	23	8	2	36
Optimization time (s)				18000	18000	18000	40	1200

Table 4: Summary of optimization results for the liver case for three points standing on Pareto front illustrated on Figure 6a

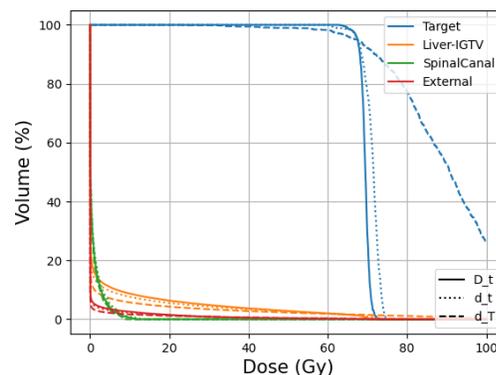
4. Discussion

We have presented the first application of **bi-criteria** optimization to the PAT treatment planning problem, which requires innovative trade-offs between clinical and delivery time goals.

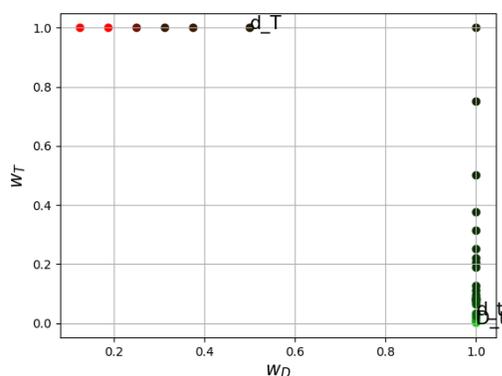
Mixed-Integer programming. Finding an optimal point on Pareto front is converted into a traditional single objective optimization problem. We choose mixed-integer linear programming as most suited formalism to tackle it, and a well known optimizer to solve it. According to our previous study (Wuyckens et al. 2022), the MIP problem statement is a proper formulation, allowing to model our full objective function and constraints in a convex form, which is a commonly admitted prerequisite to apply the BCO algorithm. Given enough computation time, solutions with optimality proofs can be obtained by using an appropriate solver (such as Gurobi) on our MIP model. Unfortunately, our problem is proven to be NP-hard (Wuyckens et al. 2022), which means that such proofs, or even the corresponding solutions, are impossible to compute in reasonable time, for realistically sized problem instances. This is the case when the dimension of the problem increases, typically for higher dose resolution or larger tumor. This has been observed in some of the Pareto plots, where several points found by the optimizer were slightly **“off the curve”**. This indicates a clear limitation of our current optimizer to give an efficient approximation of the Pareto frontier. However, looking at the results, we can say that our optimizer is close to finding and proving the optimal solutions considering the smooth shape of the Pareto fronts with a slight experimental noise that also supports



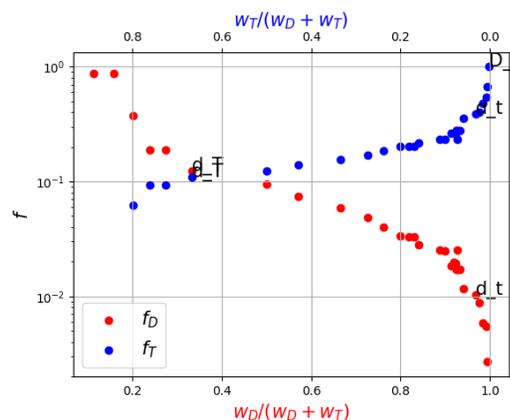
(a) Pareto front



(b) DVHs associated with 3 points on Pareto front



(c) Absolute weighting factors distribution



(d) Relative weighting factors distribution

Figure 6: BCO results obtained for the liver case. D_t is a solution where dose fidelity objective has nearly exclusive priority versus $BDT \approx$ objective while d_T just do the opposite and d_t is a more balanced combination of both objectives.

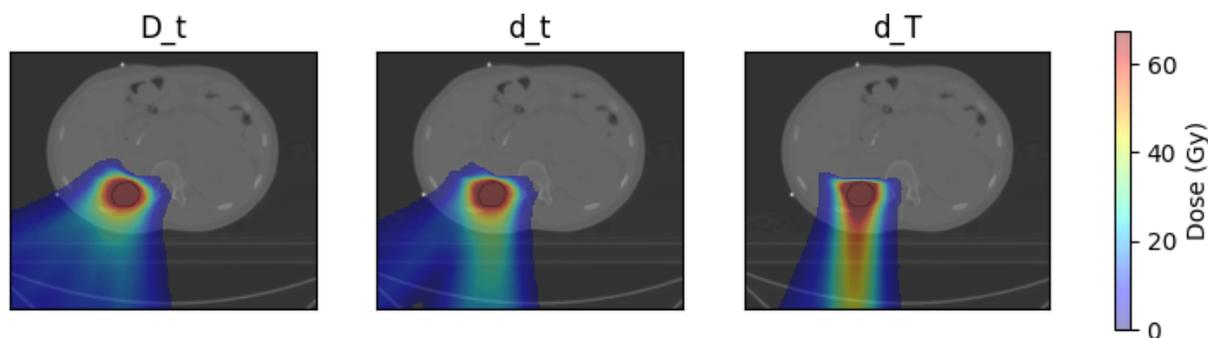


Figure 7: Liver dose distributions for three points lying on Pareto front.

the claim of a convex approximation. Over 30 optimizations for each case, about 10 are useful to represent the Pareto front. Some runs were indeed not necessary because

their final solution was already found in a previous run or eventually ended being not Pareto optimal. We suspect this behavior is induced by the discrete nature of the PAT problem. Because we are using a MIP model, the number of points in the Pareto front should be finite. The Pareto front is therefore discrete and disconnected, as observed in the results. Another benefit of our MIP formulation is its ability to interchange soft objectives with hard constraints, although this has not been exploited in the current article. As shown in Wuyckens et al.'s (2022), the ability to turn medical constraints, from soft to hard, is mandatory in order to compute solutions with strong treatment guarantees. Finally, the benchmark produced for this study including a conventional IMPT plan and a SPArc plan allows to make several statements. Given that for each case, the plan parameters had to be adjusted to bring sufficient target coverage for both IMPT and SPArc plans, it only reflects the power of degrees of freedom that is efficiently used in the MIP plans. It also demonstrates that the energy sequence optimization matters and especially for more complex cases such as small brain tumors for which, in our case study, SPArc (through its greedy heuristic) was unable to satisfy the dose constraints even with tuned plan parameters.

Anchor points choice. The weights of the anchor points we selected to start the BCO algorithm might not be very reasonable considering a large part of the Pareto optimal plans are too extreme for a physicist in real clinical practice. Indeed, many points lead to poor results in terms of dose distribution and would not be regarded at all. A less exhaustive and more economical exploration technique should therefore be devised to avoid some unnecessary optimization runs. For example, fixing hard constraints (made possible by MIP) to obtain a minimal acceptable target coverage in addition to the soft objectives might have given more clinically relevant results.

True delivery time. In this work, we only consider an approximation of the *static* beam delivery time. However, the *true* delivery time is in reality *dynamic* and includes mechanical parameters on top of the irradiation parameters such as maximum gantry velocity, acceleration, and deceleration speeds. The modeling of such function and the search for a suitable optimizer are very complex and require further research.

Time considerations. Solving each scalarized problem for a mixed-integer programming problem is very costly. For this study, each optimization has been run for 5 hours. Considering we have made 30 iterations to approximate a single Pareto front for each test case, it implies that 6 days in total were needed, which is certainly not realistically feasible for a clinical implementation at the present time. Nevertheless, this huge computation time is not exclusively attributable to the BCO layer. Underneath, PAT optimization is still recent and much progress is needed to improve the optimization efficiency. In our case, further effort could be dedicated to the MIP problem statement itself, as well as to the tuning of the optimizer parameters.

N dimension. We have considered a bi-objective formulation although in reality the fidelity objective is itself a multi-objective problem. The planners state a new objective for each target/organ clinical goal specified in the prescription. A weighting factor is therefore attached to each term and requires a fine tuning to fulfill the combination of objectives. In this work, we indicated that these factors were adjusted beforehand through an iterative manual loop. However, one could ask about the implications for subsequent bi-objective optimization. The reason behind this choice is the following: the dose quality is the most important objective in any case, as we cannot deliver a plan that does not respect the dose limits. If we find a good set of weighting factors for the target and OARs that work to obtain a good plan quality, we assume the same set will also work for the subsequent bi-objective optimization including the BDT objective given that we only seek to minimize BDT as long as the dose quality remains sufficient. We understand that the last statement is weak and might deteriorate our algorithm efficiency to solve the PAT problem since degrees of freedom are cut out when adjusting the weights beforehand or that a set of weights might be more convenient than another for a specific BDT weight but this extends outside the scope of this work. The next logical step is therefore to use multi-criteria optimization. One could use the convex hulling method presented by Craft et al. (2006) to take all the dimensions of the problem into account. For example, in the liver case, the MCO would show the trade-offs between average liver dose, high dose volumes for the other OARs, or conformality. Advanced techniques for generating Pareto fronts of MIP-specific problems should also be considered though today this is still a very active field of research. Recently, Burachik et al. (2019) presented a review of such algorithms where they ultimately proposed a new four-objective algorithm in a challenging mixed-integer programming problem.

Uncertainties. In this study, uncertainties have voluntarily been neglected in the treatment planning optimization and, consequently, no robustness evaluation was performed. However, robustness planning and evaluation are of utmost importance for proton therapy due to dose distributions sensitivity to the uncertainties of the position of the Bragg peaks (Albertini et al. 2011, Fredriksson 2012). However, the integration of robustness in proton arc therapy treatment optimization is extremely difficult and also computationally too intensive at this time, when using optimizers with no restrictive pre-selection of the energy layers, *i.e.*, using the full available set of degrees of freedom. Actually, none of the works published in the literature on proton arc optimization with no pre-selection has presented results optimized with robust settings. In our case, mixed-integer programming will require elaborated mathematical techniques in order to obtain robust treatment plans. Given a set of sampled outcomes, the L-shaped method (Laporte & Louveaux 1993) could be exploited to minimize expected deviations from the initial treatment. Monte Carlo sampling, such as the sample average approximation method (Ahmed, Shapiro & Shapiro 2002), should also be considered, although the complexity of our problem will severely limit the size of the sample required to produce estimates. Few methods exist in order to reduce the size of a sample, while preserving its repre-

sentativeness (Kahn & Marshall 1953). One could also look at the progress made with the well-known beam angle optimization (BAO) problem (Cao et al. 2022) and translate the methods to energy layer optimization (ELO) problem we have in proton arc therapy. Moreover, there is still limited data for robustness optimization or evaluation for PAT, and therefore no indication on what a good robustness evaluation strategy would possibly be used for this new modality. For this purpose, the difference of robustness between PAT and IMPT should be investigated and, based on that, tools to evaluate the robustness of proton arc plans can be developed in a clinically relevant way. Note that there is still no consensus for the robustness evaluation of a conventional IMPT plan itself, though the topic is widely recognized by the community (Sterpin et al. 2021).

MCO utility. If planners made use of an efficient MCO algorithm, it would eventually speed up planning time by a great amount and ease the decision-making process to select the plan, which represents the most desirable compromise between target coverage, organ at risk sparing, and irradiation time, where the last is critical to make efficient use of the PAT modality. While interesting algorithmic challenges remain open for PAT treatment plan optimization (reliable approximation of beam delivery time, MIP modeling) as well as for Pareto front generation in multi-objective mixed-integer programming problems, we believe PAT potential could ultimately be unlocked when combined with MCO.

5. Conclusion

PAT is a new modality that delivers the treatment over multiple proton beam directions following an arc. The PAT optimization problem imposes minimizing the beam delivery time in addition to the traditional clinical and dosimetric goals. Trade-offs between these two objectives are therefore indispensable in order to deliver dose and time efficient PAT plans. In this work, we have demonstrated the feasibility of applying a **bi-criteria** algorithm to the PAT problem statement, which could further help the treatment planners understand the different options that are actually available. The next steps as future work would be to further reduce the optimization time, e.g., by using an iterative grid size approach, implementing heuristics to shrink the searchable solution space, etc. Moreover, scaling to N dimension, with more than 2 concurrent objectives, and hence solving a clinically relevant MCO is another important milestone on the future roadmap of PAT.

Acknowledgments

S.W. is funded by the Walloon Region as part of the Arc Proton Therapy convention (Pôles Mecatech et Biowin). L.Z. is supported by a Beaumont Research Seed Grant. J.A.L. is a Research Associate with the F.R.S.-FNRS. X.D. reports honorarium from IBA and Elekta speaker Bureau, as well as research fundings from IBA, Elekta,

and RadioMed outside the work presented here. Computational resources have been provided by the supercomputing facilities of the Université catholique de Louvain (UCLouvain/CISM) and the Consortium des Équipements de Calcul Intensif en Fédération Wallonie Bruxelles (CÉCI) funded by the F.R.S.-FNRS under convention 2.5020.11.

References

- Ahmed, S., Shapiro, A. & Shapiro, E. (2002), ‘The sample average approximation method for stochastic programs with integer recourse’, *Submitted for publication* pp. 1–24.
- Albertini, F., Hug, E. B. & Lomax, A. J. (2011), ‘Is it necessary to plan with safety margins for actively scanned proton therapy?’, *Physics in Medicine and Biology* **56**(14), 4399–4413.
URL: <https://doi.org/10.1088/0031-9155/56/14/011>
- Bahr, G. K., Kereiakes, J. G., Horwitz, H., Finney, R., Galvin, J. & Goode, K. (1968), ‘The method of linear programming applied to radiation treatment planning’, *Radiology* **91**(4), 686–693. PMID: 5677503.
URL: <https://doi.org/10.1148/91.4.686>
- Bertolet, A. & Carabe, A. (2020), ‘Proton monoenergetic arc therapy (PMAT) to enhance LETd within the target’, *Physics in Medicine & Biology* **65**(16), 165006.
URL: <https://doi.org/10.1088/1361-6560/ab9455>
- Bortfeld, T. (1999), ‘Optimized planning using physical objectives and constraints’, *Seminars in Radiation Oncology* **9**(1), 20–34. Radiation Therapy Treatment Optimization.
URL: <https://www.sciencedirect.com/science/article/pii/S1053429699800526>
- Brahme, A. (1995), *Treatment Optimization Using Physical and Radiobiological Objective Functions*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 209–246.
URL: https://doi.org/10.1007/978-3-662-03107-0_1
- Breedveld, S., Craft, D., van Haveren, R. & Heijmen, B. (2019), ‘Multi-criteria optimization and decision-making in radiotherapy’, *European Journal of Operational Research* **277**(1), 1–19.
URL: <https://www.sciencedirect.com/science/article/pii/S0377221718307148>
- Burachik, R. S., Kaya, C. Y. & Rizvi, M. M. (2019), ‘Algorithms for generating pareto fronts of multi-objective integer and mixed-integer programming problems’.
URL: <https://arxiv.org/abs/1903.07041>
- Cao, W., Rocha, H., Mohan, R., Lim, G., Goudarzi, H., Ferreira, B. & Dias, J. (2022), ‘Reflections on beam configuration optimization for intensity-modulated proton therapy’, *Physics in Medicine and Biology* **67**(13). Funding Information: We thank Amy Ninetto, Scientific Editor, Research Medical Library, MD Anderson Cancer Center, for editing the manuscript. This work is partly supported by the National Cancer Institute of the National Institutes of Health through grants R03 CA256220, U19 CA021239 and Cancer Center Support (Core) Grant P30 CA016672 to the University of Texas MD Anderson Cancer Center. This work has been partly supported by the Fundação para a Ciência e a Tecnologia (FCT) under project grants UTA-EXPL/FMT/0079/2019, PTDC/CCIINF/28030/2017, UIDB/00645/2020, UIDB/05037/2020 and UIDB/00308/2020. Publisher Copyright: © 2022 Institute of Physics and Engineering in Medicine.
- Clausen, J. (2003), Branch and bound algorithms-principles and examples.
- Craft, D., Halabi, T. & Bortfeld, T. (2005), ‘Exploration of tradeoffs in intensity-modulated radiotherapy’, *Physics in Medicine and Biology* **50**(24), 5857–5868.
URL: <https://doi.org/10.1088/0031-9155/50/24/007>
- Craft, D. L., Halabi, T. F., Shih, H. A. & Bortfeld, T. R. (2006), ‘Approximating convex pareto surfaces in multiobjective radiotherapy planning’, *Medical Physics* **33**(9), 3399–3407.
URL: <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.2335486>

- Craft, D., McQuaid, D., Wala, J., Chen, W., Salari, E. & Bortfeld, T. (2012), ‘Multicriteria vmat optimization’, *Medical Physics* **39**(2), 686–696.
URL: <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.3675601>
- Ding, X., Li, X., Zhang, J. M., Kabolizadeh, P., Stevens, C. & Yan, D. (2016), ‘Spot-scanning proton arc (sparc) therapy: The first robust and delivery-efficient spot-scanning proton arc therapy’, *International Journal of Radiation Oncology*Biography*Physics* **96**(5), 1107–1116.
- Engwall, E., Battinelli, C., Wase, V., Marthin, O., Glimelius, L., Bokrantz, R., Andersson, B. & Fredriksson, A. (2022), ‘Fast robust optimization of proton PBS arc therapy plans using early energy layer selection and spot assignment’, *Physics in Medicine & Biology* **67**(6), 065010.
- Fredriksson, A. (2012), ‘A characterization of robust radiation therapy treatment planning methods—from expected value to worst case optimization’, *Medical Physics* **39**(8), 5169–5181.
URL: <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.4737113>
- Gu, W., Ruan, D., Lyu, Q., Zou, W., Dong, L. & Sheng, K. (2020), ‘A novel energy layer optimization framework for spot-scanning proton arc therapy’, *Medical Physics* **47**(5), 2072–2084.
- Gurobi Optimization, LLC (2022), ‘Gurobi Optimizer Reference Manual’.
URL: <https://www.gurobi.com>
- Kahn, H. & Marshall, A. W. (1953), ‘Methods of reducing sample size in monte carlo computations’, *Journal of the Operations Research Society of America* **1**(5), 263–278.
- Küfer, K.-H., Hamacher, H. W. & Bortfeld, T. (2000), A multicriteria optimization approach for inverse radiotherapy planning.
- Laporte, G. & Louveaux, F. V. (1993), ‘The integer l-shaped method for stochastic integer programs with complete recourse’, *Operations Research Letters* **13**(3), 133–142.
URL: <https://www.sciencedirect.com/science/article/pii/016763779390002X>
- Lawler, E. L. & Wood, D. E. (1966), ‘Branch-and-bound methods: A survey’, *Operations Research* **14**(4), 699–719.
URL: <https://doi.org/10.1287/opre.14.4.699>
- Liu, G., Zhao, L., Zheng, W., Qin, A., Yan, D., Janssens, G., Shen, J., Stevens, C., Kabolizadeh, P. & Ding, X. (2020), Develop an accurate model of spot-scanning treatment delivery time for a compact superconducting synchrocyclotron proton system.
URL: <https://w3.aapm.org/meetings/2020AM/programInfo/programAbs.php?sid=8491aid=52335>
- Marsten, R., Subramanian, R., Saltzman, M., Lustig, I. & Shanno, D. (1990), ‘Interior point methods for linear programming: Just call newton, lagrange, and fiacco and mccormick!’, *Interfaces* **20**(4), 105–116.
URL: <https://doi.org/10.1287/inte.20.4.105>
- Morrison, D. R., Jacobson, S. H., Sauppe, J. J. & Sewell, E. C. (2016), ‘Branch-and-bound algorithms: A survey of recent advances in searching, branching, and pruning’, *Discrete Optimization* **19**, 79–102.
URL: <https://www.sciencedirect.com/science/article/pii/S1572528616000062>
- Oelfke, U. & Bortfeld, T. (2001), ‘Inverse planning for photon and proton beams’, *Medical Dosimetry* **26**(2), 113–124.
URL: <https://www.sciencedirect.com/science/article/pii/S0958394701000577>
- Romeijn, H. E., Ahuja, R. K., Dempsey, J. F. & Kumar, A. (2006), ‘A New Linear Programming Approach to Radiation Therapy Treatment Planning Problems’, *Operations Research* **54**(2), 201–216.
- Shaw, E., Kline, R., Gillin, M., Souhami, L., Hirschfeld, A., Dinapoli, R. & Martin, L. (1993), ‘Radiation therapy oncology group: Radiosurgery quality assurance guidelines’, *International Journal of Radiation Oncology*Biography*Physics* **27**(5), 1231–1239.
- Souris, K., Lee, J. A. & Sterpin, E. (2016), ‘Fast multipurpose monte carlo simulation for proton therapy using multi- and many-core cpu architectures’, *Medical Physics* **43**(4), 1700–1712.
- Sterpin, E., Rivas, S. T., den Heuvel, F. V., George, B., Lee, J. A. & Souris, K. (2021), ‘Development of robustness evaluation strategies for enabling statistically consistent reporting’, *Physics in*

Medicine & Biology **66**(4), 045002.

URL: <https://doi.org/10.1088/1361-6560/abd22f>

Wuyckens, S., Saint-Guillain, M., Janssens, G., Zhao, L., Li, X., Ding, X., Sterpin, E., Lee, J. A. & Souris, K. (2022), ‘Treatment planning in arc proton therapy: Comparison of several optimization problem statements and their corresponding solvers’, *Computers in Biology and Medicine* p. 105609.

URL: <https://www.sciencedirect.com/science/article/pii/S0010482522004012>

Yu, Y., Zhang, J., Cheng, G., Schell, M. & Okunieff, P. (2000), ‘Multi-objective optimization in radiotherapy: Applications to stereotactic radiosurgery and prostate brachytherapy’, *Artificial intelligence in medicine* **19**, 39–51.

Zhang, G., Shen, H., Lin, Y., Chen, R. C., Long, Y. & Gao, H. (2022), ‘Energy layer optimization via energy matrix regularization for proton spot-scanning arc therapy’, *Medical Physics* **Accepted**.

Zhao, L., Liu, G., Chen, S., Shen, J., Zheng, W., Qin, A., Yan, D., Li, X. & Ding, X. (2022), ‘Developing an accurate model of spot-scanning treatment delivery time and sequence for a compact superconducting synchrocyclotron proton therapy system’, *Radiation Oncology* **17**(1).

URL: <https://doi.org/10.1186/s13014-022-02055-w>

Appendix A. MIP model

We represent a specific sequence of energy layers (EL) as a path in a directed graph, from a source node 0^- which has outgoing directed edges $(0^-, n)$ to every EL $n \in E$, to a sink node 0^+ which has incoming edges from every EL, passing by a network of edges that link the ELs altogether: there is an edge (m, n) for each pair of ELs such that n belongs to a strictly higher beam than m . We note E^{edges} the set of all edges that constitute the graph, whereas $E^- = E \cup \{0^-\}$, $E^+ = E \cup \{0^+\}$, $E^\pm = E \cup \{0^-, 0^+\}$ denotes the nodes.

Let x_j and e_{mn} be decision variables such that

- $\forall j \in S$: x_j determines the intensity attributed to a spot $j \in S$,
- e_{mn} are redundant variables determining whether some energy layer (EL) switch happens: $\forall m, n \in E^\pm$: $e_{mn} = 1$ if $(m, n) \in E^{\text{edges}}$ and layer $n \in E^+$ is activated directly after layer $m \in E^-$, otherwise $e_{mn} = 0$.

From Wuyckens et al. (2022), the proton arc therapy treatment optimization problem for finding a single optimal point on the Pareto front can be modelled as a weighted sum using a mixed-integer two-index flow formulation

$$\min_{x_1, \dots, x_{|S|}} \quad F(\mathbf{x}) := w_D f(\mathbf{d}) + w_T \text{BDT}^\approx(\mathbf{x}) \quad (\text{A.1})$$

$$\text{s.t. :} \quad \sum_{n \in E^+} e_{0^-n} = \sum_{n \in E^-} e_{n0^+} = 1 \quad (\text{A.2})$$

$$\forall m \in E \quad \sum_{l \in E^-} e_{lm} = \sum_{n \in E^+} e_{mn} \quad (\text{A.3})$$

$$\forall m \in E, \forall j \in S_m \quad x_j \leq I \sum_{l \in E^-} e_{lm} \quad (\text{A.4})$$

$$\forall j \in S \quad 0 \leq x_j \leq I \quad (\text{A.5})$$

$$\forall (m, n) \in E^{\text{edges}} \quad e_{mn} \in \{0, 1\} \quad (\text{A.6})$$

where the fidelity term $f(\mathbf{d})$ is represented as follows:

$$f(\mathbf{d}) = \sum_{k \in \text{ROI}} u_k f_k(\mathbf{d}) \quad (\text{A.7})$$

$$\text{s.t.} \quad f_k(\mathbf{d}) = \|(d_i - p_i)_+\|_1 \quad (\text{A.8})$$

$$\mathbf{d} = \mathcal{A}\mathbf{x} \quad (\text{A.9})$$

and the approximation of the irradiation time $\text{BDT}^\approx(\mathbf{x})$ is modeled as:

$$\text{BDT}^\approx(\mathbf{x}) = \sum_{(m,n) \in E^{\text{edges}}} c_{mn} e_{mn}$$

$$\text{s.t.} \quad \forall (m,n) \in E^{\text{edges}} \quad e_{mn} \in \{0,1\}$$

$$c_{mn} = \begin{cases} 6.5 & \text{for upwards energy switching} \\ 1.6 & \text{for downwards energy switching} \\ 1.2 & \text{for constant energy} \end{cases}$$

Appendix B. The BCO Pareto approximation algorithm

This appendix provides details on the BCO algorithm through a simple example as well as a the full description of the equations derived for the Pareto approximation construction.

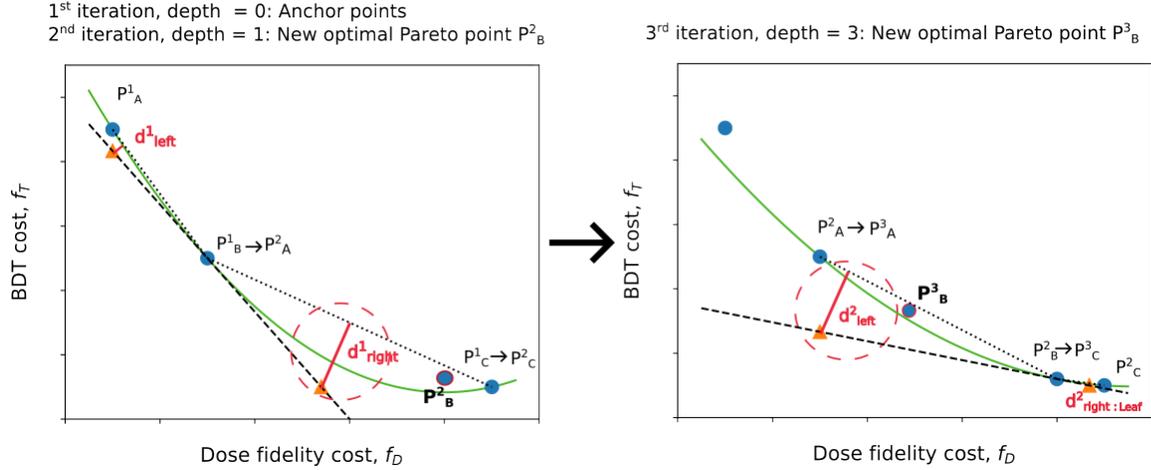


Figure B1: Pareto front geometric construction. Left (iteration 1 & 2): Constructing depth 1 on the right hand side branch of the tree depicted in Figure 1. Right (iteration 3): Constructing depth 2. Blue: Pareto optimal points (the red encircled ones are selected for next iteration). Green: upper Pareto quadratic approximation. Distal points (orange) and their distance (red) to segment lines.

Following the path depicted in the binary tree search (dashed nodes in Figure 1), we give hereinafter a detailed description of the different steps within the iterations.

Iteration 1. The first iteration starts from the parent node where we compute the so-called anchor points $\mathbf{P}_A^1, \mathbf{P}_C^1$ representing the individual minima of each objective. In the present study, these points are obtained by optimizing the global objective function with individual weights $\mathbf{w} = (w_D, w_T)$ for the dose fidelity objective f_D and the BDT \approx objective f_T . These weights are set to $\mathbf{w}_A = (1, 0.001)$, $\mathbf{w}_C = (0.001, 1)$, respectively, to avoid too poor quality solutions in terms of dose coverage. Next, a third (interior) point \mathbf{P}_B^1 is obtained with equal weights, $\mathbf{w}_B = (1, 1)$. From these three points, an initial upper Pareto quadratic approximation is obtained. Using the bounding segment lines between these points and the tangent line at the interior point, initial lower distal points and their distance to segment lines $\mathbf{P}_A\mathbf{P}_B$ and $\mathbf{P}_B\mathbf{P}_C$ (d_{left}^1 and d_{right}^1 , respectively) can be computed by solving the equation system: (see equations B.3 and B.4 in the formulas below). This step is illustrated geometrically on Figure B1 (left). At this stage, the parent node has now two children nodes corresponding to each distance.

Iteration 2. In the next iteration $i = 2$, we pick the tree node associated with the currently largest d^{i-1} distance. In our example, d_{right}^1 (resulting from P_B^1 and P_C^1) is picked, the new weights to run to obtain a new Pareto optimal point (P_B^2) are obtained by averaging the weights of P_B^1 and P_C^1 . The same geometric construction as in iteration 1 is carried out to compute the segment lines d_{left}^2 and d_{right}^2 .

Iteration 3. As d_{right}^2 is lower than a predefined tolerance threshold, we then cut the current branch (*leaf* node in tree). The third iteration thus continues with the tree node associated with d_{left}^2 . It generates a new point P_B^3 (Figure B1, right) and so on.

These iterations (summarized in Algorithm 1) are repeated until all distances d^i fall below minimal distance tolerance, or the user-specified tree depth I_{max} is reached.

In order to mitigate the bias introduced by objectives having different magnitudes, a dimension scaling on the individual objective functions is performed whenever a new point \mathbf{P}_B^i is computed, still following Craft et al.'s (2006) method.

Formulas

The anchor points for input upper Pareto approximation are $\mathbf{P}_A = (x_A, y_A)$ for $f(\mathbf{d})$ and $\mathbf{P}_C = (x_C, y_C)$ for BDT $\approx(\mathbf{x})$, the single interior point is $\mathbf{P}_B = (x_B, y_B)$. Then the upper Pareto quadratic approximation is

$$y = \frac{(x - x_C)(x - x_B)}{(x_A - x_C)(x_A - x_B)}y_1 + \frac{(x - x_A)(x - x_B)}{(x_C - x_A)(x_C - x_B)}y_C + \frac{(x - x_A)(x - x_C)}{(x_B - x_A)(x_B - x_C)}y_B$$

The slope of its tangent line is

$$y' = \frac{(x - x_C) + (x - x_B)}{(x_A - x_C)(x_A - x_B)}y_A + \frac{(x - x_A) + (x - x_B)}{(x_C - x_A)(x_C - x_B)}y_C + \frac{(x - x_A) + (x - x_C)}{(x_B - x_A)(x_B - x_C)}y_B$$

The tangent line at the interior point has the slope

$$y'(x_B) = \frac{x_B - x_C}{(x_A - x_C)(x_A - x_B)}y_A + \frac{x_B - x_A}{(x_C - x_A)(x_C - x_B)}y_C + \frac{(x_B - x_A) + (x_B - x_C)}{(x_B - x_A)(x_B - x_C)}y_B$$

The tangent line at the interior point is

$$y - y_B = y'(x_B)(x - x_B)$$

The initial lower distal point (LDP) to facet $\mathbf{P}_A\mathbf{P}_B$ is

$$\mathbf{I}_{\text{left}} = (x_A, y_B + y'(x_B)(x_A - x_B)) \quad (\text{B.1})$$

The initial LDP to facet $\mathbf{P}_C\mathbf{P}_B$ is

$$\mathbf{I}_{\text{right}} = (x_B + \frac{y_C - y_B}{y'(x_B)}, y_C) \quad (\text{B.2})$$

The equation of facet $\mathbf{P}_A\mathbf{P}_B$ is

$$y - y_A = \frac{y_B - y_A}{x_B - x_A}(x - x_A)$$

that is, in standard form

$$(y_B - y_A)x - (x_B - x_A)y + y_A(x_B - x_A) - x_A(y_B - y_A) = 0$$

The distance from LDP \mathbf{I}_{left} to facet $\mathbf{P}_A\mathbf{P}_B$ is

$$d_{\text{left}} = \frac{|(y_B - y_A)x_A - (x_B - x_A)[y_B - y'(x_B)(x_B - x_A)] + y_A(x_B - x_A) - x_A(y_B - y_A)|}{\sqrt{(x_B - x_A)^2 + (y_B - y_A)^2}} \quad (\text{B.3})$$

Similarly, the equation of facet $\mathbf{P}_C\mathbf{P}_B$ is

$$(y_B - y_C)x - (x_B - x_C)y + y_C(x_B - x_C) - x_C(y_B - y_C) = 0$$

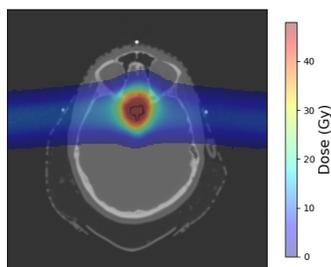
The distance from LDP $\mathbf{I}_{\text{right}}$ to facet $\mathbf{P}_C\mathbf{P}_B$ is

$$d_{\text{right}} = \frac{|(y_B - y_C)[x_B + \frac{y_C - y_B}{y'(x_B)}] - (x_B - x_C)y_C + y_C(x_B - x_C) - x_C(y_B - y_C)|}{\sqrt{(x_B - x_C)^2 + (y_B - y_C)^2}} \quad (\text{B.4})$$

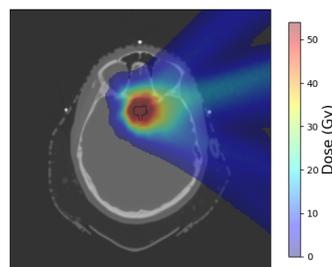
Appendix C. **Benchmark**

For each disease site, IMPT and SPArc dose distributions are shown in Figure C1. The DVHs are provided in Figure C2. To give an appreciation of the MIP plan quality compared to IMPT and SPArc, a pastel shade band was added on top of the DVH curves, with a thickness proportional to the range of solutions between the MIP plans D.t and d.T.

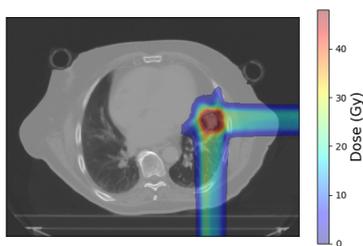
In Section 3, we have argued that the plan parameters for the IMPT and SPArc plans had to be fine-tuned from those used in the MIP plans, in order to reach a sufficient target coverage. We provide additional material to support this assertion. In Table C1, original plan parameters (used for the MIP plans) and tuned parameters are listed. Table C2 reports the target coverage obtained for each plan and each patient, as well as the initial number of spots at plan creation and thus available to the optimizer.



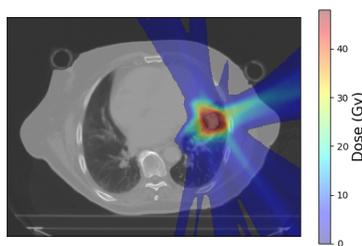
(a) IMPT - Brain



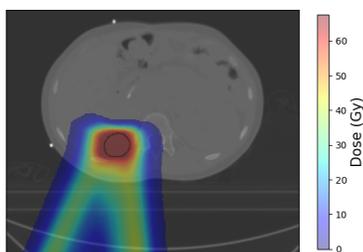
(b) SPArc - Brain



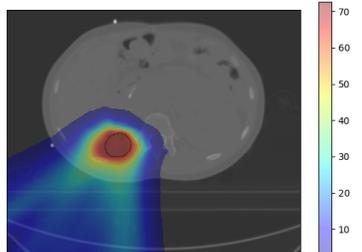
(c) IMPT - Lung



(d) SPArc - Lung

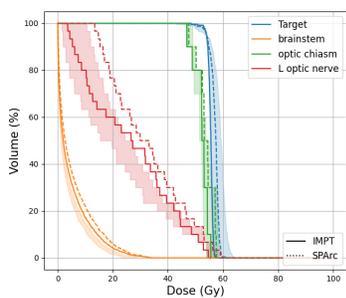


(e) IMPT - Liver

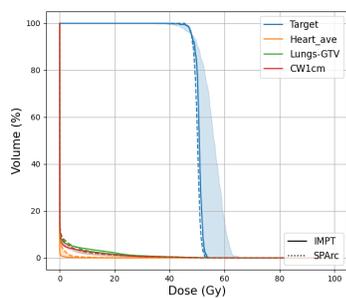


(f) SPArc - Liver

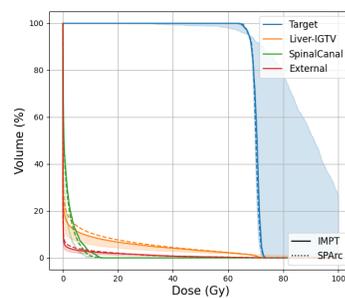
Figure C1: Optimal dose distributions for 3 patients



(a) Brain



(b) Lung



(c) Liver

Figure C2: DVH comparison between IMPT (solid line), SPArc (dashed line) and MIP-BCO solutions (filled area)

Plan parameters	Disease site	Original	Tuned
Spot spacing [mm]	Brain	7	5
	Lung		7
	Liver		5
Layer spacing [mm]	Brain	4	3
	Lung		4
	Liver		3
Target margin [mm]	Brain	6	7
	Lung		
	Liver		

Table C1: Plan parameters for spot placement

	Modality	IMPT		SPArc	
	Plan parameters	Original	Tuned	Original	Tuned
D5-D98 (Gy)	Brain	7.5	3.4	10.9	6.3
	Lung	5.1		6.9	4.6
	Liver	18.5	5.8	11.8	5.7
# spots	Brain	459	1352	475	614
	Lung	312		472	550
	Liver	880	2481	1217	2698

Table C2: Target coverage comparison of plans using original (MIP) and tuned plan parameters