

Treatment planning in arc proton therapy: comparison of several optimization problem statements and their corresponding solvers

Sophie Wuyckens^{a,*}, Michael Saint-Guillain^b, Guillaume Janssens^c, Lewei Zhao^d, Xiaoqiang Li^d,
Xuanfeng Ding^d, Edmond Sterpin^{a,e}, John A. Lee^a, Kevin Souris^a

^a*UCLouvain, Molecular Imaging, Radiotherapy and Oncology (MIRO), Brussels, Belgium*

^b*UCLouvain, ICTEAM, Louvain-La-Neuve, Belgium*

^c*Ion Beam Applications SA, Louvain-La-Neuve, Belgium*

^d*Department of Radiation Oncology, Beaumont Health, Royal Oak, MI, USA*

^e*KULeuven, Department of Oncology, Leuven, Belgium*

Abstract

Arc proton therapy (ArcPT) is a new modality in cancer treatments. It delivers the proton beams following a sequence of irradiation angles while the gantry is continuously rotating around the patient. Compared to conventional proton treatments (intensity modulated proton therapy, IMPT), the number of beams is significantly increased bringing new degrees of freedom that leads to potentially better cancer care. However, the optimization of such treatment plans becomes more complex and several alternative statements of the problem can be considered and compared in order to solve the ArcPT problem. Three such problem statements, distinct in their mathematical formulation and properties, are investigated and applied to solving the ArcPT optimization problem. They make use of (i) **fast iterative shrinkage-thresholding algorithm** (FISTA), (ii) local search (LS) and (iii) mixed-integer programming (MIP). The treatment plans obtained with those methods are compared among them, but also with IMPT and an existing state-of-the-art method: Spot-Scanning Proton Arc (SPArc). MIP stands out at low scale problems both in terms of dose quality and time delivery efficiency. FISTA shows high dose quality but experiences difficulty to optimize the energy sequence while LS is mostly the antagonist. This detailed study describes independent approaches to solve the ArcPT problem and depending on the clinical case, one should be cautiously picked rather than the other. This paper gives the first formal definition of the problem at stake, as well as a first reference benchmark. Finally, empirical conclusions are drawn, based on realistic assumptions.

Keywords: Cancer, Proton therapy, Algorithms, Optimization, Planning, Treatment, Arc proton

*Corresponding author.

Full postal address: Claude Bernard Avenue Hippocrate 54, boîte B1.54.07, 1200, Woluwe-Saint-Lambert, Belgium

Email address: sophie.wuyckens@uclouvain.be (Sophie Wuyckens)

1. Introduction

At this time, arc proton therapy (ArcPT) has not yet any commercial available software solution to optimize treatment plans. Moreover, due to the increased degrees of freedom and complexity of optimization along the arc trajectory, generating an ArcPT treatment plan efficiently is one of the hindrances for a future clinical implementation. A treatment plan normally includes all parameters that allow the calculation of radiation dose to be delivered to the patient. A prototype of the ArcPT delivery machine, built by the IBA company, exists and is being studied at the Beaumont clinic in the United States [17] (Fig. 2). The finalized product is still under development by IBA. In this study we have developed and investigated three optimization techniques, namely, fast iterative shrinkage-thresholding algorithm (FISTA), local search (LS), and mixed-integer programming (MIP). We compared them with the existing state-of-the-art iterative SPArc optimization framework [8].

Contributions. The main contributions of this work are: **(i)** We introduce the first reference benchmarks for this problem, generated through several processing steps performed on a *realistic* medical image set. The benchmarks are accessible from Open Source Dataverse [33] and can be used freely for testing. **(ii)** Three new optimization algorithms to generate ArcPT treatments plans are proposed. **(iii)** Whereas most existing studies showcase a single plan optimization method in several different disease sites, this is the first comprehensive and quantitative comparison across several existing and new ArcPT optimization techniques.

2. Background, challenge and related work

Background. Radiotherapy (RT) treats cancer with ionizing radiation, mostly x-ray photons. However, protons offer unique physical properties by stopping in the tumor and releasing most of their energy at the end of their track, known as the Bragg Peak [32]. The interest in proton therapy (PT) delivery has therefore increased. This modality has a clinical potential over RT because radiation dose deposited beyond the tumor can be avoided while at the same time the proximal dose remains low, thereby reducing the radiation toxicity in healthy tissues. Nowadays, state-of-the-art PT treatments are delivered through intensity modulated proton therapy (IMPT) using pencil beam scanning (PBS)[12]. Typically made up of 2 to 4 *fixed-angle* beams, as illustrated in Fig. 1 (left), IMPT delivers each beam

with multiple smaller intensity modulated beams [22], called spots, each reaching a specific location of the target. Moreover, to irradiate the target in depth, multiple energy layers per beam angle (also referred to as incidence angle or irradiation angle) are needed. Therefore, an IMPT treatment is delivered spot-by-spot within a same energy layer, layer-by-layer within a beam, and then beam-by-beam.

ArcPT builds on top of IMPT in the sense that, instead of delivering the treatment with a few *fixed-angle* beams, it uses a gantry that *rotates continuously* around the patient with a variable speed and delivers the treatment following a predefined discrete series of irradiation angles, called the angular control points. An arc is hence defined by a start, stop, and step angles and by specifying the gantry rotation direction. This is depicted in Fig. 3. ArcPT is expected to increase treatment conformity and robustness [28] and to achieve better clinical goals in terms of tumor coverage and organs at risk sparing [27].

Challenge. When the number of beam angles is increased (*e.g.* from 3 to 46) to form an arc and the amount of energy layers in each beam of the arc is kept as numerous as in the IMPT treatment (Fig. 1, left), usually in the order of tens, a so-called arc multi-field plan is obtained (Fig. 1, center). Unfortunately, this high number of beams (and energy layers) results in a significant increase in the complexity of the optimization problem, computationally demanding in time and memory. Moreover, although conducting to potentially better treatments, the plan obtained on Fig. 1 (center) would, above all, not be delivered in practice due to the huge delivery time required to deliver each set of energy layers in each beam which affects the patient’s comfort and directly relates to the number of patients that can be treated in a day. Indeed, the PT system spends a large part of the delivery time to switch the energy between layers, namely, the energy layer switching time (ELST). Therefore, the challenge of ArcPT plan optimization lies in selecting as few energy layers as possible, ideally only one for each beam direction to reduce the ELST and obtain deliverable ArcPT plans (Fig. 1, right). Moreover, due to the magnetic hysteresis effect, the PT device takes more time to increase the energy than decreasing it, which means that the energy switch-downs should be favored over the switch-ups, to further reduce the ELST. The key to obtain such sparsity in energy is thus the selection and the sequence of the energy layers while keeping sight of the dosimetric objectives to get a fully-optimized treatment plan.

State-of-the-art. Today, only a few research groups have investigated this research topic and published methods of optimizing ArcPT plans. For instance, RaySearch (RaySearch Laboratories AB, Stock-

holm, Sweden), has recently published a proof-of-concept method of energy layer and spot selection, coined ELSA (Early Layer and Spot Assignment), prior to spot weight optimization [9]. They were able to produce high-quality ArcPT plans in a fast optimization process at the cost of losing degrees of freedom in the solution space, caused by these well-separated steps for ELSA and spot weight optimization. Furthermore, an oncology group from UCLA has published some results on a parsimonious algorithm named ELO-SPAT [10], minimizing an objective function made of several terms. The terms control the selection and sequence of the energy layers, therefore including a delivery time criterion. However, they resort to a non-convex term that can degrade the algorithm performances. Additionally, a group from Beaumont clinic in the US, developed the Spot-Scanning Proton Arc (SPArc) algorithm [8][18], an iterative greedy approach based on robust optimization (what-if analysis). At the moment, this is the most popular and practical method for its simplicity and efficiency. It has been applied to various disease sites and has shown remarkable results [21][4][6][20][7][16]. The major drawback in the SPArc is the greediness of its iterative approach, with an optimization process that can spend much time in refining a solution that might eventually prove to be suboptimal in terms of plan quality. In this paper, three new algorithms that can tackle the ArcPT challenges are presented and compared with each other, as well as with existing state-of-the-art methods, in terms of plan quality, energy layer sparsity, delivery efficiency and, lastly, optimization and dose calculation speed.

3. Proposed Benchmark

This section describes the different steps that were performed to generate the ArcPT treatment plans based on a primary image set acquired with Computed Tomography (CT). For this study, a public, open-access, computed tomography (CT) phantom image set collected at MAASTRO Clinic (Maastricht, NL)[11] in DICOM format was elected. In the field of medical imaging, a phantom is a specially designed object or synthetic image created to experimentally evaluate and tune the treatment or imaging workflow. Here, the phantom simulates the abdomen of a small adult.

Generalities about proton therapy Treatment Planning. The design of plans follows a well established procedure made up of three typical stages applicable to all PBS-based PT treatments. They are illustrated in Fig. 4 and described hereafter:

- (1) **Spot placement.** This first stage determines, for each beam angle chosen beforehand, a

list of energy layers (EL) to cover the target in depth (the higher the proton energy the deeper they penetrate in the target). Then for each EL, a grid of irradiation points (namely spots) is generated. At this step, the plan is thus already defined by the list of spots, with their x and y positions, their energy, and angular incidence.

- (2) **Beamlet calculation.** The patient geometry is discretized in small three dimensional volume elements called *voxels*. With a particular spot (*i.e.*, within a specific energy layer) is associated a so-called *beamlet*. As depicted in Fig. 4, a beamlet describes how the irradiation (elementary dose) from the spot is distributed inside the patient, that is, over the voxels. These beamlets are obtained by simulating the underlying proton physics, using a dose calculation algorithm based on a Monte Carlo engine, MCsquare, developed in-house [31]. This type of algorithm is valued for its accuracy although it can be slow.
- (3) **Spot weight optimization.** This is the core of this work, the spot weights (also referred to as spot intensities) optimization for ArcPT modality. The total dose delivered to the patient can be computed as the sum of all beamlets weighted by the intensity of the corresponding spots (See Section 4 Eq. 2). Therefore, we only need to optimize the weights of these spots so that the total dose map meets the clinical objectives and constraints. Compared to conventional IMPT plans, new degrees of freedom (the numerous angular positions) and new constraints (treatment delivery time) arise and make the ArcPT optimization problem more complex. New plan optimization methods to determine ArcPT spot weight are implemented in *OpenTPS*, an in-house treatment planning system (TPS).

ArcPT benchmark. In our test sets, the beam angles chosen in step (1) were sampled along a quarter of arc (90°) set with a start angle (90°), stop angle (0°), and a patient couch angle (0°), yielding a counter-clock wise gantry rotation. We decided to produce arc plans with 5 different step angles 1° , 2° , 3° , 5° , and 10° , yielding plans with 91, 46, 31, 19, and 10 beams, respectively. The spot placement was then performed for each beam direction [26] by generating multiple hexagonal grids stacked in depth to fully cover the 3D target volume with spots separated by a pre-defined layer spacing and spot spacing. Prior to computing the adequate spot intensities (*i.e.*, solving the optimization problem), the beamlets were calculated (step (2)) based on the spot placement using the MCsquare dose engine [31]. Benchmarks of different sizes and accuracy were

eventually obtained by computing these beamlets under 5 different image resolutions (from 1 to 5mm) using a resampling operation of the DICOM files. A summary of the benchmarks with the time and memory allocated to produce the beamlet files is reported in Table 1. All the datasets we produced are publicly available from Open Source Dataverse [33].

4. Notations and Problem Definition

4.1. Notations

We can now introduce notations that will be useful within our optimization framework. In this paper, the notation used by Gu *et al.* [10] has been slightly modified for the sake of clarity and ease of reading.

- B is the sequence of beams forming an arc, defining the angular control points of the gantry and spaced by a pre-defined step angle; the sequence is ordered by irradiation angle, following the gantry rotation;
- E_b is the sequence of candidate energy layers for beam $b \in B$, ordered by descending energy value;
- S is the set of all the spots defined in the plan;
- S_e is the set of candidate spots for layer $e \in E$;
- $V = T \cup R$ are the sets of target (T) and OAR (R) voxels defined in the patient geometry;
- \mathcal{A} is the beamlet matrix, containing $|S| \times |V|$ elementary dose values (where $|\cdot|$ is the cardinality operator);
- \mathbf{x} is the column vector of spot weights, containing $|S|$ values.

Therefore, the dose map is computed formally as:

$$\mathbf{d} = \mathcal{A}\mathbf{x} = \sum_{j \in S} \mathcal{A}_j x_j , \quad (1)$$

$$d_i = \sum_{j \in S} \mathcal{A}_{ij} x_j, \quad \mathcal{A}_{ij} \geq 0 , \quad (2)$$

where $\mathbf{d} \in \mathbb{R}_{\geq 0}^{|V|}$ is the vectorized dose and d_i is the elementary dose delivered to voxel i . Eq. 1 encodes the physics of the treatment, *i.e.*, the description of the dose contribution for each spot, as a linear combination of the beamlets \mathcal{A} (input data) and the spot intensities \mathbf{x} (decision variables).

For practical reasons, we will mostly use a grouped version of \mathbf{x} where each vector component represents the spot intensities within beam b , \mathbf{x}_b . The latter is itself a vector where each component represents the spot weights within beam b and layer e , \mathbf{x}_e . We also use a compact representation of \mathbf{x} and \mathbf{x}_b , \mathbf{y} and \mathbf{y}_b that are vectors with each element representing the sum of weights in each beam and each layer, respectively:

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_b \\ \vdots \\ \mathbf{x}_{b'} \end{bmatrix}, \mathbf{x}_b = \begin{bmatrix} \mathbf{x}_e \\ \vdots \\ \mathbf{x}_{e'} \end{bmatrix}, \mathbf{y} = \begin{bmatrix} \mathbf{y}_b \\ \vdots \\ \mathbf{y}_{b'} \end{bmatrix}, \mathbf{y}_b = \begin{bmatrix} y_e \\ \vdots \\ y_{e'} \end{bmatrix}, y_e = \sum_{j \in S_e} x_j, \quad (3)$$

where $b, b' \in B$ and $e, e' \in E_b$ and eq. 3 can be written as a matrix multiplication in a more compact way:

$$\mathbf{y} = W \mathbf{x}$$

with W , the summation matrix along the spot dimension.

4.2. Problem statement, constraints, and evaluation

The problem faced by the physicists in ArcPT treatment planning can be ideally formulated by Program (4-8). The solution is a vector of spot weights \mathbf{x} that minimizes the beam delivery time $\text{BDT}(\mathbf{x})$ subject to medical constraints, *i.e.* target volumes (TV) coverage and organs at risk (OAR) sparing with minimal, maximal and/or mean prescription and dose limits respectively,

$$\min_{\mathbf{x}} \text{BDT}(\mathbf{x}) \quad (4)$$

$$\text{s.t.} \quad \forall k \in \text{TV}, \forall i \in T_k : \quad p_{\min T_k} \leq d_i \leq p_{\max T_k} \quad (5)$$

$$\forall k \in \text{OAR}, \forall i \in R_k : \quad d_i \leq p_{\max R_k} \quad (6)$$

$$\forall k \in \text{OAR} : \quad \frac{1}{|R_k|} \sum_{i \in R_k} d_i \leq p_{\max \text{mean} R_k} \quad (7)$$

$$\mathbf{x} \in \mathbb{R}_{\geq 0}^S \quad (8)$$

with $|T_k|$ and $|R_k|$, the total number of voxels associated with TV k and OAR k and p_{maxR_k} stating a maximal irradiation for any voxel i in a given OAR k , for example the maximal dose in the spinal cord must not exceed 45 Gy [15].

However, plan optimization methods in the current literature fail at strictly considering clinical constraints (5)-(7), leading to OARs being irradiated over the acceptable maximal (resp. average) dose constraint (6) (resp. constraint (7)) if, for example, they are in direct contact with the tumor. Instead, existing methods, usually based on gradient-descent (continuous) optimization, approximate them with soft constraints $f(\mathbf{d})$ that penalize an overall (alternative) objective function:

$$\min_{\mathbf{x}} \text{BDT}(\mathbf{x}) + f(\mathbf{d}) \quad (9)$$

$$\text{s.t.} \quad \mathbf{d} = \mathcal{A}\mathbf{x} \quad (10)$$

$$\mathbf{x} \in \mathbb{R}_{\geq 0}^{|S|} \quad (11)$$

where $f(\mathbf{d})$ is a penalty function evaluating the quality of the simulated treatment dose *w.r.t.* constraints (5)-(7), formulated as a weighted sum of multiple clinical goals depending on the priorities of the regions of interests (ROI, composed of both TVs and OAR volumes):

$$f(\mathbf{d}) = \sum_{k \in ROI} w_k f_k(\mathbf{d}) \quad (12)$$

Depending on the clinical objective, $f_k(\mathbf{d})$ is typically formulated as a squared l_2 -norm of the difference between the dose $\mathcal{A}\mathbf{x}$ and prescription \mathbf{p} . In practice, multiple clinical goals are encoded using weighted soft objectives to constrain the minimum, maximum, or mean doses. This way two slightly offset one-sided quadratic functions can be used to constraint the dose within a band to which we can give either low or large importance weight. The total cost function can then explicitly written as:

$$f(\mathbf{d}) = \sum_k \frac{w_k}{|T_k|} \sum_{i \text{ in } T_k} \min(0, d_i - p_{maxT_k})^2 \quad (13)$$

$$+ \sum_k \frac{w_k}{|T_k|} \sum_{i \text{ in } T_k} \max(0, d_i - p_{minT_k})^2 \quad (14)$$

$$+ \sum_k w_k \max(0, \frac{1}{|R_k|} \sum_{i \text{ in } R_k} d_i - p_{maxmeannR_k})^2 \quad (15)$$

$$+ \dots \quad (16)$$

Program (9-11) now properly describes the ArcPT problem and shows a real trade-off between the smooth part dealing with the clinical goals $f(\mathbf{d})$ and the non-smooth part minimizing the beam delivery time $\text{BDT}(\mathbf{x})$. Because a real delivery time model will be certainly not convex and hardly differentiable, we choose to optimize the energy sequence in the absence of a better model of the BDT. Indeed, the BDT is predominantly driven by the energy layer switching time (ELST) in ArcPT therefore we use the ELST as a surrogate to the BDT in our optimization problem. In order to simplify the model, we assume that it takes 5.5 seconds to switch the energy up versus 0.6 seconds to switch it down [19]. In this study, we thus assume the ELST to be the sum of all individual switching times between the active energy layers (layers with non-null spot weights) at each angular control point along the arc.

Evaluation. To assess a plan and score its quality, the physicists compute dose-volume histograms (DVH) and the associated DVH metrics for both TVs and OARs. A DVH (see for example Fig. 10) shows a decumulative distribution of dose within some given ROI. A point on the curve represents the percentage of the considered ROI that receives a dose less than or equal to the dose value corresponding. As we want to irradiate the tumor volume, the ideal curve is a step from 100% to 0% at the prescribed dose to avoid under- and over-dosage. On the other hand, we want to spare healthy organs. This translates as a ideal step from 100% to 0% at null dose. Obviously, in practice, actual DVH curves generally look like smoothed steps. In addition to DVH curves, DVH metrics further summarize the dose map information into a few scalar indicator. For instance, D_{95} represents the minimal dose received by at least 95% of the considered volume (usually a target volume).

5. State-of-the-art methods

5.1. SPArc

Developed by Ding *et al.* [8], SPArc is a greedy heuristic planning optimization algorithm that uses two nested iterative loops to optimize the dose distribution while increasing the number of control points to make it denser until a desired arc delivery sampling frequency is obtained. It starts with a sparse arc, *i.e.*, with very few irradiating directions. Subsequently, the nested loop optimization algorithm can start. The spot weights are optimized for the current beam-layer configuration. Four steps follow:

1. Control point resampling: each beam is split into two or more new irradiation directions

2. Energy layer redistribution between the new sub-control points
3. Energy layer filtration to get rid of low weighting energy layers

At each iteration of optimization, steps (1) to (3) are carried out until the angle sampling frequency reaches the user-defined parameter. At this stage, SPARC ensures acceptable delivery time by designing plans with 1 to 3 active layers per control point. However, the energy switching pattern is still arbitrary. SPARC has therefore been improved [18] to re-distribute the energy layers in a more efficient way, namely, in descending order to mitigate the impact from energy-switch upward on the delivery time.

Because of its simplicity and proven potential in several studies [4][6][20][16] this algorithm is appreciated in the community. However, its heuristic and greedy character makes it less efficient in the treatment plan optimization process. Additionally, the spot weights and the energy sequence are not optimized jointly but as separate steps and this may result in sub-optimal delivery efficiency and plan quality. Finally, improvement of plan quality or delivery efficiency of the final treatment plan may require a reset of the initial sampling process.

5.2. ELO-SPAT

Gu et al. developed the algorithm ELO-SPAT [10]. It solves the optimization problem by including a regularization strategy, to address the ArcPT non-smooth problem stated in Eq. 9-11. In particular, it uses the proximal gradient method (PGM)[1], a popular technique to solve non-smooth optimization problems for image reconstruction. However, traditional PGM methods tend to be quite slow on convergence, so an accelerated proximal gradient scheme is implemented in the computation. The algorithm chosen is the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [2]. The ELO-SPAT framework is formulated with an objective function made of several terms to meet all ArcPT constraints. Based on notations introduced in Section 4, Eq. 17 gives an intuition on the ELO-SPAT objective function:

$$\arg \min_{\mathbf{x}} \|\mathcal{A}\mathbf{x} - \mathbf{p}\|^2 + \sum_{b \in B} \sum_{e \in E_b} \alpha_e \|\mathbf{x}_e\|_2^{1/2} - \beta \sum_{b \in B} \log \left(\sum_{e \in E_b} y_e \right) + \gamma \sum_{b \in B} h(\tilde{\mathbf{y}}_b) \quad \text{s.t. } \mathbf{x} \geq 0, \mathbf{y} = W\mathbf{x} \quad (17)$$

The first term, namely, the dose fidelity term, encodes the objectives of the treatment. It is a soft-constraint version of constraints (5)-(7), where the OAR volumes are naturally constrained with a zero dose. It implicitly contains a weighted sum of several objectives (Eq. 12), each being a squared l_2 -norm between the prescription to tumor (and dose limits to OARs) and voxels of the total dose distribution estimate. The second term is a group sparsity regularization defined by a non-convex $l_{2,1/2}$ -norm that induces a sparsity at the level of the energy layers, *i.e.* to obtain as few energy layers per beam as possible resulting in a reduced ELST. This term is weighted by the hyperparameter α_e and is responsible to turn off non-interesting layers (low weight in layer). The third term represented by a log barrier behaves as a counter-acting term to the sparsity term in order to distribute the selected layers across the whole gantry rotating range, a heuristic for obtaining better arc-shaped treatments. Finally, in the fourth term, a regularization function is used to sequence the energies, in a way that tends to minimize the delivery time (Eq. 4). This function asymmetrically penalizes energy switching low-to-high harder than high-to-low, using a one-sided quadratic cost function h that acts on the energy differences between adjacent beams where the layer with the maximal spot intensity \tilde{y}_b has been selected beforehand. In terms of workflow, the algorithm starts from all possible EL's candidates, select as few as possible and then encourage energy switch from high to low.

This algorithm raised a lot of attention because it is written in a very well defined mathematical framework. ELO-SPAT converges towards the global or local minimum (due to non-convex term) faster compared to SPARC from a mathematical point of view. However, there is still room for improvement, in particular, in the formulation of some objectives. Firstly using a non-convex group sparsity term together with FISTA is not proper given that FISTA has been designed to work with convex objective functions. Next, the log barrier is a really strong term for the intended goal since it uses a continuous approximation of a penalty that should be null on a large domain. The penalty starts to increase too early and might shift the optimum where this is not necessary. Finally, for the last term, they use specially designed discrete gradient operators that need to be updated after every iteration which is again not ideal for FISTA convex optimization. We worked to improve the formulation of the ArcPT cost function based on ELO-SPAT.

6. Proposed plan optimization methods

In this section, we present three new, fundamentally different plan optimization methods. One is an improved variant of ELO-SPAT, thus based on gradient descent. The second method is a local search

approach, based on simple neighborhood operators and a Simulated Annealing meta-heuristic. The third method is a Mixed-Integer Programming formulation. As discussed later in Section 9, these three methods come with significantly different mathematical properties in terms of optimality, treatment guarantees, and computational effort.

6.1. FISTArcPT

FISTArcPT (eq. 18) is a modified version of the ELO-SPAT objective function formulation:

$$\arg \min_{\mathbf{x}} f(\mathbf{d}) + \lambda \sum_{b \in B} \sum_{e \in E_b} \alpha_e \|\mathbf{x}_e\|_2^1 + \sum_{b \in B} I_C^b(\mathbf{x}) + \gamma \sum_{b \in B^-} \text{LeakyReLU}(D_b^B \text{LSE}(\mathbf{z}_b \cdot \tanh(a \cdot \mathbf{y}_b)) , \quad \text{s.t. } \mathbf{x} \geq 0, \mathbf{y} = W\mathbf{x} , \quad (18)$$

where B^- represents the set B of beams, without the last beam of the sequence B . Each term of Eq. 18 is described by order of appearance, from left to right, below in details.

1. Dose fidelity

The dose fidelity term $f(\mathbf{d})$ packs multiple objectives into a sum of weighted soft objectives as described in Eq. 14-16. This term is the core of the treatment plan optimization and occurs in any RT or PT problem formulation.

2. Group sparsity

For the second term of Eq. 18, we use the group LASSO [30], that is, the sum of squares of coefficients belonging to the same group and, in this case, the weights belonging to the same energy layer. It is embodied in a weighted $l_{2,1}$ -norm, contrasting with the non-convex $l_{2,1/2}$ -norm used in the original ELO-SPAT method. The main reason is that the $l_{2,1}$ is convex (Fig. 5), which fulfills the requirements of the FISTA algorithm. Moreover, the proximal operator is simpler to implement and still efficient for our purpose, the deactivation of layers. The inner l_2 -norm of $\|\mathbf{x}_e\|_2^1$ enforces selection of all the weights within an energy layer while the outer l_1 -norm promotes sparsity in the number of selected energy layers. It is controlled by two parameters, namely, the constant λ set at the beginning of the optimization and α_e that is updated every ten iterations and defined as:

$$\alpha_e = \begin{cases} \frac{1}{y_e} = \frac{1}{\sum_{j \in S_e} x_j} & \text{if } y_e \geq 0, \\ 1 & \text{otherwise} \end{cases} . \quad (19)$$

With proper tuning of λ , this term selects fewer layers out of all candidate layers.

3. Indicator function

The group sparsity term, if used solely, can result in a solution with some beams having multiple layers activated and some with none, which reduces the arc geometry benefits. To counteract this effect, we implemented an indicator function I_C defined in Eq. 20, as an alternative to the log barrier in ELO-SPAT:

$$I_C(x) = \begin{cases} 0 & \text{if } x \in C, \\ \infty & \text{if } x \notin C \end{cases}, \quad (20)$$

where the admissible domain C is defined as

$$C = \{y_e : \sum_{e \in E_b} y_e \geq \epsilon\}.$$

In practice, the function operates if the sum of weights in a beam is less than the tunable threshold, ϵ . If so, it projects the weights of active layers belonging to the beam under consideration such that the sum of weights in these layers is exactly ϵ . If no layer is active in the beam, it projects directly the weights of the layer with the minimal gradient on ϵ . This is a smooth way to ensure that no zero intensity beam remains and that a certain homogeneity in the arc is kept. By combining the group sparsity term with the indicator function, on average one layer per beam is selected. However, the energy sequence is not yet optimized.

4. Energy sequencing

To encourage energy switches from high to low and, therefore, to minimize the treatment delivery time, we devised a differentiable regularization term controlled by the weighting parameter γ . A hyperbolic tangent is first used to map the sum of intensities of each layer between 0 and 1 with a a scaling factor tuned for that purpose. Next, this vector is multiplied element-wise by z_b , the vector of energies corresponding to each layer. A log-sum-exp operator is then used as a soft maximum to select the most eligible energy layer in each beam. After this, a discrete gradient operator D_b^B is applied to each beam. For example, $D_b^B \mathbf{y}_b = \mathbf{y}_{b+1} - \mathbf{y}_b$, *i.e.*, the vector of the intensity difference between the two beams considered. Finally a leaky ReLU function is used to

asymmetrically penalize the energy switching pattern in the following way:

$$f(x) = \begin{cases} 0.01x & \text{if } x < 0 \text{ (switch-down)} \\ x & \text{else (switch-up).} \end{cases} \quad (21)$$

Optimizer. The cost function is non-differentiable due to $l_{2,1}$ -norm and indicator function. For this reason, we decided to use the FISTA optimizer, as in the original ELO-SPAT method, which converges at an optimal rate of $\mathcal{O}(1/k^2)$ [2]. The FISTA optimizer engine was implemented in *OpenTPS* with the Python package PyUNLocBoX [5] that is based on the original FISTA algorithm [2]. The pseudo code of FISTArcPT is described in Algorithm 1 and mainly consists in the forward step performing the gradient descent (dose fidelity and energy sequencing) followed by the backward step achieved through the proximal operator of the non smooth part of the objective function, *i.e.*, the group sparsity term.

Algorithm 1 Pseudo code for FISTArcPT algorithm

```

1: Input: Initial weights  $x_0$ , step  $\alpha \geq 0$ ,  $\lambda$ ,  $\epsilon$ ,  $\gamma$ 
2: for  $i = 1$  to  $N$  do
3:    $x_k = \text{update\_sol}(x_k)$            // FISTA acceleration
4:    $\alpha = \text{update\_step}(x_k)$        // Backtracking
5:    $x_{k+1} \leftarrow x_k - \alpha \cdot \nabla f(x_k)$  // Forward step
6:    $x_{k+1} \leftarrow I_{\geq 0}(x_{k+1})$  // Positive projection
7:    $x_{k+1} \leftarrow \text{prox}_{\lambda \|\mathbf{x}\|_2^1}(x_{k+1})$  // Backward step
8:    $x_{k+1} \leftarrow I_C(x_{k+1})$  // Indicator II step
9: end for

```

Convexity. As mentioned in Subsection 5.2, FISTA has been developed to solve convex optimization problems. By analyzing each term of the FISTArcPT problem statement (Eq. 18), we can check the convexity of our objective function. Firstly, the fidelity term involves soft objectives in the form of half-quadratics that are convex functions. By definition, the sum of convex functions is also convex, so is our total dose fidelity term. The second term is the group sparsity $l_{2,1}$ -norm purposely chosen for its convex nature, illustrated in Fig. 5. The third term, the indicator function, is convex by observation. Fig. 6 exposes in a simplified schema the convexity of the set C .

We were not able to analytically determine the convexity of the last objective, the energy sequencing (ES) function. Instead, a Monte Carlo simulation suffices to prove the non-convexity of the ES function. Given that the energy sequencing term is a sum of functions over the different beams of the plan, the non-convexity of only one term in the sum needed to be proved. For some beam b , we evaluate the function $\text{LeakyReLU}(D_b \text{LSE}(z_b \cdot \tanh(a \cdot \mathbf{y}_b)))$. For one evaluation of this function, at minimum four

variables were needed due to the discrete nature of the gradient that requires two beams (b and $b + 1$) to operate and also to the fact that \mathbf{y}_b is a vector containing at least two variables. Jensen’s inequality [23] was used as convexity criterion requiring the function to be evaluated twice. The Monte Carlo simulation therefore needs $(N^4)^2 = N^8$ input variables. Running the simulation during one hour that generates $N = 1000$ points per iteration adding up to 23673 iterations, we computed the proportion of non-convex points to be about 35.7%. Unfortunately, a convex form of the energy sequencing term could not eventually be found. This conclusion leads to the fact that the optimization problem as stated in Eq. 1 is non-convex. Therefore, the optimization algorithm FISTA is likely to converge to local minima and suffers in its performances.

6.2. Local Search optimizer

Alternatively to the gradient-based methods, we propose a simpler method based on the local search (LS) algorithm, described in Algorithm 2. It implements a Simulated Annealing [13] meta-heuristic for approximating the optimal solution x^* , minimizing

$$f(x) = f(\mathbf{d}) + \alpha \text{ sparsity} + \gamma \text{ ELST} . \quad (22)$$

Starting from an initial incumbent (*i.e.*, current) solution x , Algorithm 2 iteratively modifies it by using a set of five neighborhood operators. At each iteration, it randomly chooses a solution x' in the current neighborhood $\mathcal{N}(x)$ (line 4) and either accepts it and resets the distance counter d to zero (line 6) or rejects it and increments d (line 11). Whether a candidate solution is accepted as the new incumbent solution or not depends on a predefined *acceptance criterion* function, described below, which implements the so-called *Simulated Annealing*. A restart strategy is also implemented in line 14: whenever the current incumbent solution becomes too distant (in terms of local changes) from the last best solution x^* , the search simply restart from x^* . At the end, the algorithm simply returns the best solution x^* encountered so far, which is stored at line 8.

Initial solution. If not specified, the initial solution is constructed by simply assigning a zero value to every spots. Otherwise, any possible value assignment of the spots can be considered as an initial incumbent solution for the LS algorithm.

Stopping criterion. The stopping criterion designed here simply depends on the computational time budget that the user allocates to the algorithm. Once the allotted time elapsed, the criterion is met.

Algorithm 2 Local Search algorithm

Let $x = \{x_s : s \in S\}$ be an *initial solution*.
2: Set $x^* \leftarrow x$; Set the distance counter d to 0
while some *stopping criterion* is not met **do**
4: Select a solution x' uniformly at random in $\mathcal{N}(x)$
 if some *acceptance criterion* is met on x' **then**
6: set x to x' and d to 1;
 if $f(x') < f(x^*)$ **then**
8: $x^* \leftarrow x'$
 end if
10: **else**
 $d \leftarrow d + 1$
12: **end if**
 if $d = d^{\max}$ **then**
14: restart: set $x \leftarrow x^*$
 end if
16: **end while**
return x^* , the best solution found

Notice that alternative stopping criteria may be devised, based on convergence indicators, for example: relative quality of the last improvements, delays between previous improvements, etc.

Neighborhood operators. We consider five neighborhood operators in $\mathcal{N}(x)$: **i**) select a spot at random and set its value to zero; **ii**) select a spot at random and assign a random positive value in $]0, M]$, where M is the highest possible intensity; **iii**) select an activated layer at random and set all its spots to value zero; **iv**) select a layer at random and reassign random positive values to all its spots; **v**) find the target voxel for which the total dose is the farthest from the prescription, select one of its contributing spots at random and assign it the (non-negative) value that minimizes the dose difference from that voxel to the prescription. At each iteration, the local search algorithm simply selects an operator at random. We notice that classical improvement avenues for such neighborhood based optimization method include an adaptive mechanism that selects the current operator probabilistically, according to its success rate observed so far, preferring operators that contributed at generating relatively many interesting solutions, that is, that passed the acceptance criterion. Such a mechanism is, however, typically implemented whenever the neighborhood operators include more complex and computationally demanding operations (*e.g.*, large neighborhood search, [24]).

Incremental computation. An important property of the proposed set of neighborhood operators is that given a current solution x , any neighboring solution $x' \in \mathcal{N}(x)$ is actually constructed incrementally from x . In particular, the new objective value $f(x')$ is also incrementally computed provided $f(x)$. In

all our operators (\mathbf{i} to \mathbf{v}), both computations of x' and $f(x')$ are done in $\mathcal{O}(|S|)$ linear time, relatively to the number $|S|$ of spots. Incremental computation and evaluation is in fact a necessary property of any efficient local search approach. On the contrary, every iteration of gradient based methods such as SPArc, ELO-SPAT or FISTArcPT, involve the heavy matrix-vector product Ax of worst-case $\mathcal{O}(VS^2)$ complexity.

Acceptance criterion. We use a Simulated Annealing (SA) acceptance criterion [13]. Improved solutions are always accepted, while degraded solutions are accepted with a probability that depends on the degradation magnitude and a temperature parameter, *i.e.* the probability of accepting x' is $e^{-\frac{1-f(x)/f(x')}{T}}$. The temperature T is updated by a *cooling factor* $0 < f_T < 1$ at each iteration of Algorithm 2: $T \leftarrow f_T \cdot T$. During the search process, T gradually evolves from an initial temperature T_{init} to nearly zero. We also implement a *restart strategy*: each time the current temperature T decreases below a fixed limit T_{min} , T is reset to T_{init} . In all experiments, SA parameters were set to $T_{\text{init}} = 5, T_{\text{min}} = 10^{-6}$, and $f_T = 0.8$.

6.3. Mixed-Integer Programming Formulation

An exact alternative to approximate optimization methods as presented above is to use a MIP formulation and an appropriate solver, allowing a complete search of the solution space to be conducted in order to prove optimality. Furthermore, the treatment requirements can be guaranteed, as hard constraints, instead of using soft constraints with minimized violations, as it is often the case with gradient-based methods, hence approximating the initial (*i.e.*, ideal) treatment problem stated in (4)-(8).

Energy sequencing. 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. Remark that the digraph resulting from E^{edges} allows for valid energy sequencing only. In particular, a unique path from 0^- to 0^+ enforces that *at most one* EL is activated in each beam; otherwise, the path would contain forks. Finally, a non-negative value c_{mn} gives the cost of using edge $(m, n) \in E^{\text{edges}}$ in the path, that is, the cost of switching from layer m to layer n , as illustrated in Fig. 7.

6.3.1. Linear model

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 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 $e_{mn} = 0$.

Then, the arc proton therapy treatment optimization problem can be modelled as a mixed-integer two-index flow formulation

$$\min_{x_1, \dots, x_{|S|}} \sum_{i \in R, j \in S} \mathcal{A}_{ij} x_j \quad (23)$$

$$\text{s.t.: } \forall i \in T \quad p_{\min T} \leq \sum_{j \in S} \mathcal{A}_{ij} x_j \leq p_{\max T} \quad (24)$$

$$\sum_{n \in E^+} e_{0^-n} = \sum_{n \in E^-} e_{n0^+} = 1 \quad (25)$$

$$\forall m \in E \quad \sum_{l \in E^-} e_{lm} = \sum_{n \in E^+} e_{mn} \quad (26)$$

$$\sum_{(m,n) \in E^{\text{edges}}} c_{mn} e_{mn} \leq U \quad (27)$$

$$\forall m \in E, \forall j \in S_m \quad x_j \leq I \sum_{l \in E^-} e_{lm} \quad (28)$$

$$\forall j \in S \quad 0 \leq x_j \leq I \quad (29)$$

$$\forall (m, n) \in E^{\text{edges}} \quad e_{mn} \in \{0, 1\} \quad (30)$$

The objective function (23) minimizes the total dose on the OAR voxels, *provided* (24) each target voxel $i \in T$ receives its prescribed treatment, here in the range $[p_{\min T}, p_{\max T}]$, *and* (27) the total energy sequencing cost does not exceed a predefined bound U . In order to stick with the previous methods, we assume $c_{mn} = 1$ when m is of lower energy than n , 0 otherwise, so that the total summation simply gives the number of EL switch-ups. The set of constraints (25)-(26) are classical flow conservation constraints, ensuring that the selected edges from E^{edges} form a unique path (without fork) from source 0^- to sink 0^+ .

The relation between each energy layer and its associated spots is defined by constraints (28). A positive constant I represents the highest possible intensity assigned to any spot. In fact, a spot j for

which the energy layer m is not activated (*i.e.*, is not part of the flow) must necessarily be assigned energy $x_j = 0$. Finally, the aforementioned decision variables are defined in (29)-(30), where we notice that contrarily to x 's, the flow variables e_{mn} are binary integers.

Benefits and lessons from the MIP problem statement. In contrast to the previous methods, the exact MIP statement provides *guarantees* on the optimized treatment plan that will fulfill the constraints. In fact, any feasible solution for constraints (24-30) guarantees fulfilling the treatment requirements (5) on the target voxels (here that every target voxels receives a dose between p_{minT} and p_{maxT}). Otherwise, alternative statements can easily be formulated in order to enforce (6) and/or (7), while minimizing the energy sequencing cost. Furthermore, any optimal solution for program (23-30) guarantees the best possible treatment. For instance, one can set $U = 5$ and, therefore, determine whether a treatment that satisfies the $[p_{minT}, p_{maxT}]$ target dose prescription bounds actually exists when limiting to 5 EL switch-ups and what is its side effect cost in terms of OAR irradiation. Therefore, varying the EL bound U eventually provides a Pareto curve, describing the possible compromises between treatment quality and duration, for a particular use case. Finally, from a mathematical point of view, the problem statement as a linear model tells us that the energy sequencing can be represented by a convex function, not only for 0-1 switch up counts, but also for any possible cost functions to measure duration, as a simple pre-computed c_{mn} table.

6.3.2. Computational complexity

The inherent complexity of the ArcPT treatment problem, stated in (4)-(8), has not been addressed yet in the literature. Is the problem *NP-hard*? It could be that the problem is in fact polynomial, but remains intractable at this time, due to the typically huge size of the beamlet matrix \mathcal{A} . This question will likely remain open for now, although we propose a preliminary analysis, where we draw the first basic complexity boundaries by determining two variants of our problem: a special case that is polynomial and a generalization that we prove to be NP-complete.

Energy sequencing is hard. Unlike linear programming, which is known to be solvable in polynomial time, (*e.g.*, using interior point methods [25]), integer and mixed-integer programming are naturally NP-hard in general. It is easy to see that program (23)-(30), without energy sequencing, that is, when

removing constraints (25)-(28) and (30), leads to a linear program:

$$\begin{array}{ll}
\min_{x_1, \dots, x_{|S|}} & \sum_{i \in R, j \in S} \mathcal{A}_{ij} x_j \\
\text{s.t.:} & \forall i \in T \quad p_{\min T} \leq \sum_{j \in S} \mathcal{A}_{ij} x_j \leq p_{\max T} \\
& \forall j \in S \quad 0 \leq x_j \leq I \text{ ,}
\end{array}$$

which is solved in polynomial time. Therefore, if the overall treatment problem is NP-hard, the complicated part of the problem actually lies in the energy sequencing. By setting a value U that is large enough to never be reached by the summation (27), that is, *the problem becomes polynomial if one does not care about the beam delivery duration*.

NP-hardness: the case of the lower bounded energy sequencing variant. A direct generalization of the problem can be obtained by considering a lower bound L on the energy sequencing time, in addition to the upper bound U , hence replacing (27) with

$$L \leq \sum_{(m,n) \in E^{\text{edges}}} c_{mn} e_{mn} \leq U \text{ .}$$

In order to prove the hardness of this generalization, we show how to achieve the polynomial reduction from the well-known NP-complete subset sum problem (SSP): given a multiset K and a target value Y , decide whether some subset of the items of K sums exactly to Y . Let us construct our EL directed graph (that of Fig. 7), such that there is one column per item $k \in K$, only one node in each column (representing the value of k), and finally c_{mk} equals the value of the target node k , whatever the source node m . Clearly, a path in the constructed directed graph represents a particular subset selection with its associated total value. Any decision algorithm for our generalized treatment problem then decides the SSP, when $L = Y = U$. Our generalized problem is then NP-complete, as any given solution can be verified in polynomial time, (*i.e.*, the problem is in NP).

6.3.3. Large Neighborhood Search.

NP-hard or not, preliminary experiments showed that the ArcPT treatment optimization problem remains complicated to solve in practice. In fact, the disadvantage of problem statement (23)-(30), common to all complete optimization methods, is that it hardly scales up to realistically-sized problems. In this section, we propose a basic workaround, by exploiting a well-known technique: large

neighborhood search (LNS).

LNS [24] relieves the optimization process by considering, iteratively, sub-parts of the entire optimization problem. It works as follows. Given a current (sub-optimal) solution x , select a subset x_i, \dots, x_j of decision variables and fix their current values in x . This results in a sub-problem in which part of the variables are fixed, whereas the remaining variables could be optimized. By solving this sub-problem to optimality, one obtains a new current solution x' , which is necessarily as good as x . Starting from an initial solution x_{init} , this process can be repeated as long as needed.

LNS can be seen as a variant of our local search method defined in Algorithm 2, in which the only neighborhood operator \mathcal{N} consists in this variable-selection and optimization process. Because \mathcal{N} may not deteriorate the current solution x , $\mathcal{N}(x)$ is always accepted. LNS is therefore a "hill-climbing" heuristic method.

Variable selection. As the energy layers form a partition of the decision variables, it seems natural that the variable decision mechanism follows these energy layers. At each iteration of the LNS and given the current solution x , the optimization is then performed by restricting the treatment to use only a subset of the energy layers (ELs). This subset is composed of all the ELs that are currently activated in x , as well as a percentage (5% in our experiments) of randomly selected ELs. Any EL that does not belong to the selection is necessarily not activated in current x , and therefore all the associated decision variables are set to zero in the MIP model for the upcoming optimization.

Initial solution and bi-objective optimization. Finding an initial feasible solution, which satisfies constraints (24-30), is already really challenging. Instead, it would be easier to start from an empty solution, thus no activated EL, and iteratively improve the solution by activating ELs. Some iterations would then activate some ELs, whilst deactivating some of those ELs that were previously activated (this is why currently activated ELs are always part of the variable selection), and so on, hence continuously improving the current solution through the LNS iterations. The workaround we propose is then to relax constraints (24) so that an empty initial solution becomes feasible:

$$\forall i \in T : \sum_{j \in S} \mathcal{A}_{ij} x_j \leq p_{\max T}.$$

At each LNS iteration, the current MIP model (*i.e.*, having only subset of the decision variable available for optimization) is then optimized based on the following lexicographical bi-objective function

$$\max_{x_1, \dots, x_{|S|}} \sum_{i \in T, j \in S} \mathcal{A}_{ij} x_j, \quad \min_{x_1, \dots, x_{|S|}} \sum_{i \in R, j \in S} \mathcal{A}_{ij} x_j$$

, which maximizes the total treatment on the target voxels T first and minimizes the irradiation on OAR voxels R second. In other words, a soft constraint substitutes for the hard constraint on the lower bound of the target voxels' prescribed dose.

LNS: pros and cons. The direct advantage of the LNS approach is that each iteration actually ends up with a solution. In fact, for large instances, the complete MIP approach may require hours (or even not terminate) before providing a feasible, even sub-optimal, solution. By sacrificing the optimality guarantee, *i.e.*, by renouncing the model completeness, the MIP model can still be exploited to construct and improve solutions, yet sub-optimal, scaling to larger problem instances. However, in addition to optimality, with the relaxation of lower bounding constraints in Eq. 24, part of the treatment guarantees are lost with the LNS approach.

7. Comparison strategy

Plan optimization methods. From the proposed benchmark (see Section 3), ArcPT plans have been generated and optimized, using the three algorithms proposed in Section 6. To serve as first reference point, IMPT plans (fixed-angle beams) were also produced. Another reference point was given by SPArc plans (described in Section 5.1), obtained with a Python implementation of the original SPArc algorithm in our in-house TPS. All these plan optimization methods were empirically compared to each others.

Key performance indicators (KPIs). The KPIs used for this study are divided into four main categories:

1. The first KPIs in importance relate to the planned dose quality: the dose volume histograms (DVHs), re-normalized to D50 (dose \leftarrow dose \cdot $p/D50$) for an easier comparison. Quantitative KPIs are also extracted from the DVH, such as the absolute difference between D95 and D5, written as $|\Delta(D95, D5)|$, characterizing the uniformity of dose distribution within the target volume, as well as the conformity index (CI_{RTOG} [29]) that describes how tightly the prescription dose is conforming to the target.

2. The next KPIs relate to the time delivery efficiency of the plan, which depends on the sparsity of energy layers and the number of energy switches upwards or downwards, further summarized in the ELST metric used as a surrogate of the delivery time and given in the tables.
3. The third KPIs are the final numbers of spots (nonzero spot weights) and layers (*activated* layers with non-null intensity), which are provided as proxies for the degrees of freedom available to the spot weight optimizer.
4. The Last KPIs are the optimization times, which were also reported in the comparison tables. Although all the algorithms ran on the same machine, execution times should be regarded carefully because each algorithm makes use of different programming languages. FISTArcPT is written in *Python3*, whereas the local search and MIP were implemented in C++. For the last two, the optimization time depends on the time limit set in the input parameters. Furthermore, whenever the MIP found the optimal solution, it would necessarily stop before reaching the time limit.

Benchmarks. For the patient data, we are using the benchmark defined in Section 3 with clinically acceptable dose resolutions of 2 mm and 1 mm. In addition, to assess the quality of the different plans, the algorithms were executed on two target sizes: 1.47 cm^3 (*small target*) and the same but dilated using 1 cm margin (*expanded target*). The initial degrees of freedom available to the spot weight optimizer are reported in Table 2 for both modalities, *i.e.*, IMPT and ArcPT (across all methods). ArcPT plans (FISTArcPT, LS, MIP) were initialized with a pre-defined arc set between 0° and 90° with a counter-clock rotation and a step angle set to 2° . The SPArc plans were initialized with three beams (77.0° , 45.0° and 13.0°) with a maximum number of splits set to 4 so that the same arc design would be obtained. The IMPT plans were designed by initializing 3 beams at equal angular distances, at 0° , 45° and 90° , respectively.

Experimental settings. The proposed methods come with many hyper-parameters. They were tuned in a preliminary study (A) and they remained fixed in the next experiments. Notice that a convergence report of each proposed algorithm is also provided in B, as well as a description of the algorithm behaviour, for various problem scales in C. The dose prescription to the target is set to 60 Gy and no objectives nor constraints are set on the OARs because they are not included in the analysis. The maximum number of iterations, set to 4000, was used as stopping criterion for all the FISTArcPT plans. IMPT plans were optimized by applying a SciPy implementation of L-BFGS [3] in OpenTPS to a beamlet decomposition of the dose. SPArc also relies on the L-BFGS algorithm for each step of

optimization. The stopping criteria used for SciPy L-BFGS optimizations are the maximum number of iterations, set to 100, or whenever the convergence was acceptable with a relative error in function values not exceeding 0.001.

8. Results

The optimization methods presented in Sections 5-6 are intrinsically different but they all aim at designing ArcPT plans. This section demonstrates the proof of concept of each proposed method and also identifies the reasons for which one outperforms some others for specific key performance indicators.

Results on the small target. The different final dose distributions obtained with each algorithm are shown in Fig. 8 (2 mm resolution) and Fig. 9 (1 mm resolution), as well as the DVH comparison in Fig. 10. The KPIs can be compared in Table 3. In general, all the algorithms perform well, showing results that compare with the reference, that is, the IMPT treatment. FISTA struggles the hardest to reduce the number of switch-ups and to obtain a good layer sparsity with at most one energy layer active in each beam. For the lowest resolution instance (2mm), MIP performed the best and was able to reduce the number of switch-ups to zero, eventually yielding the smallest ELST, while still covering the target homogeneously. For arc plans with 1 mm of resolution, the dose conformity to the target is poor, especially for those obtained with the LS method. Hot spots appear indeed in the dose distributions (Fig. 9), because no objective in the cost function prevented the optimizer from depositing dose in any locations outside the target. Anyway, SPARC seems to give the best performances at high resolution with a very good homogeneity of the dose in the target and an acceptable time for energy layer switching.

Results on the expanded target. Dilating the target by 1 cm happens to complicate the problem considerably and the algorithms faced more difficulties to solve it, as can be observed in Table 4 and in the DVH plotted in Fig. 13, especially when the CT resolution is thinner (1 mm). Despite a good ELST, the DVH obtained with the LS algorithm is off the 5% tolerance around the prescription with a large D5. Similarly, MIP can deal with the 2-mm resolution data set but failed to produce a decent solution for the 1 mm resolution with major deviations observed in its DVH. FISTA struggles again to reduce the ELST compared to the other optimizers. In terms of dose maps (Fig. 11), the solution obtained with FISTA displays aggregated beams on the dose map despite the indicator function meant

to address this kind of issue. Moreover, the energy layers selected by each algorithm are visible on the dose distribution plots. This could be explained by the fact that the layer spacing chosen is too large or there are not enough beams to cover the energy space smoothly. Surprisingly, IMPT outperforms arc plans. Possible explanations could be that the degrees of freedom (DOF) available in the arc modality are not exploited properly by the various optimizers. As MIP starts from a solution with null spot weights and solves the problem progressively by selecting a subset of the layers (LNS approach, see Section 6.3.3), it struggles to activate spots and layers resulting in a very small final number of activated spots and layers. The same observation holds for the LS method, even though it is caused by the neighborhoods operators that (de)activate layers and spots at random. Finally, FISTA and SPARC seem to use the available DOFs in the arc better than LS or MIP, although they cannot outperform the IMPT plans in terms of dose quality. A reason could be that the optimizer needs to run longer in order to deal with these many DOFs and obtain a better solution eventually.

9. Discussion

We presented three new methods of optimizing treatment plans for arc proton therapy. Theoretical analysis and experimental evidences revealed the strengths and weaknesses of each method.

Firstly, FISTArcPT, which extends a classical gradient descent algorithm with regularization terms to control the number and the sequence of energy layers in each beam belonging to the arc, is the ideal method in the sense that it fully benefits from the arc geometry and its degrees of freedom and finds locally optimal solution with rapid convergence of objective function values. This method was able to produce good quality plans in terms of target homogeneity and conformity. However, the sensitivity to the hyper-parameters in the objective function and the need to tune them finely on case-by-case basis makes it inconvenient, especially considering the potential future use in clinic. Also, the inter-dependence between those parameters makes the algorithm biased and could lead to rapid local minima, without any means to get out of them, especially because of the definition of the shrinkage group parameter α_e , that depends on the intensity of each layer. In addition, because the objective function is not formulated to fulfill all the FISTA assumptions as defined by Beck *et al.* [2], this convex optimization iterative approach could underperform and also be trapped in local minima. Recall that we found, in Section 6, that the energy sequencing term of the objective function is actually not convex. Aside from those theoretical properties, the matrix-vector product $\mathcal{A}x$ involved in the dose fidelity term is in practice the heaviest operation in the algorithm. It is clear through

the examples that the larger the involved matrices are, with increases in resolution and/or target size, the longer the optimization time will be, and this linearly. Lastly, in addition to give an approximate solution (no treatment guarantees with soft objectives), FISTA is primarily an incomplete method, like all gradient-descent-like algorithms, because it does not consider the entirety of the solutions space and there is therefore no optimality guarantee (unless the objective function is convex, which is not our case).

Secondly, a local search approach with simulated annealing has been tested. In contrast to FISTArcPT, LS uses a randomized approach and will therefore probably never yield the same solution twice. Also, the simulated annealing meta-heuristics allows for escaping from local minima. LS was able to give solutions with high delivery efficiency (in terms of ELST) but is mostly discarded due to poor dose quality, especially for large-scale datasets. Regarding the optimization time, defined by the user time limit, it is less impacted by the size of the problem compared to FISTArcPT. In fact, in LS the computationally demanding matrix-vector product (in the gradient step of FISTArcPT) is carried out incrementally through the neighborhood operators. Besides, those neighbourhood operators remain very simple, and they could even be improved to get better results for larger-scale problems. Another point to consider is that any kind of objectives and/or constraints, including more realistic ones, such as the BDT, can be easily implemented with such algorithms, without any concerns in terms of convexity or similar properties. However, it is noteworthy that the local search algorithm is a heuristic method with no optimality guarantee and, like FISTA, it yields an approximate solution with no actual guarantee of treatment feasibility or acceptability. In any case, if the dose quality does not improve in the future upgrades of the algorithm, it cannot be used for treatment planning. Nevertheless, LS could be used as a fast pre-optimization step to obtain an initial solution with an arc configuration, such as a specific number of upwards energy switching which would then ease the search for the main, subsequent algorithm.

Thirdly, Mixed Integer Programming is a method with good potential: given enough time, it ultimately finds the optimal solution to the problem at hand, thereby guaranteeing the best possible treatment. New objectives and constraints can be easily implemented with this problem statement as long as they remain linear. With MIP, the user can, for example, set the number of switch-ups allowed in the arc at the beginning of the optimization, as well as the range of dose to the target voxels, at the expense of losing degrees of freedom of the arc method, like SPArc. However, this type of algorithm, even though giving excellent solutions for small data sets, can rapidly struggle hard once

the problem reaches a certain size, requiring hours before providing a feasible, sometimes sub-optimal solution if only it ever terminates. Several attempts have been made to address this problem scale issue. In particular, we have exploited a well-known technique: large neighborhood search (LNS) at the cost of losing part of the treatment guarantees. Another idea could be a multi-resolution approach, with a first optimization at a very low resolution to obtain a good result quickly and to input it as an initialization for a subsequent optimization at higher resolution and so on until the desired dose resolution is obtained.

Energy sequencing. The common point of all the three proposed problem statements and associated solvers (mathematically speaking, completely different from each other) is the difficulty to optimize the energy sequence as observed in the results. Stating the problem with continuous functions and using gradient-based methods such as FISTA [2], the convexity of the objective function is important to guarantee optimal performances of the solver. In the context of ArcPT, devising such functions is complex because of the discrete nature of arc therapy, leading to a delivery time estimator that is neither convex nor differentiable. Even though we simplified the real estimator by optimizing the energy sequence as a proxy of the BDT, we were not able to define a convex formulation that can be the Achilles heel of our optimization. In contrast, MIP allowed us to formulate a convex estimator for the energy sequencing, at the expense, however, of introducing integrality constraints, which complicate the problem considerably. Regarding the LS algorithm, we assume that, if the energy sequencing was not in the objectives, the cost function would be convex and a solution would be more easily found and this, even without the simulated annealing meta-heuristic. The reason is that the energy sequencing term is the cause of the local minima in which we are trapped. Yet, when we will implement a real model of the delivery time, the local search algorithm combined with the simulated annealing could become the most suitable method of beam delivery time optimization.

SPArc. Compared to SPArc, FISTArcPT was able to give similar dose quality in all cases. Only MIP could outperform SPArc in the small target cases, which indicated its high potential for the ArcPT problem. SPArc provided a reproducible plan solution (though not necessarily optimal) in both small and expanded target with different resolutions. Compared to clinical IMPT, SPArc is an advanced IMPT optimization framework that implements a greedy, iterative re-sampling process. It is noteworthy, though, that the limitations of SPArc (or other similar methods) could only appear, from a theoretical standpoint, for more complex cases with OAR sparing for example where the heuristic

greediness of SPArc and its restriction to a limited number of degrees of freedom could lead to under-performance compared to the proposed solution methods. It is therefore still difficult to demonstrate the superiority of the proposed methods with this simple test case. Yet, all the reported results should be analyzed cautiously since the benchmarks are not obtained from real patients.

Treatment uncertainties. No uncertainties have been taken into account neither in the optimization process nor in the treatment evaluation. In a clinical environment, the treatment plan has to be robust against various type of uncertainties before it can be delivered to the patient. They are usually divided into two main categories, namely, aiming errors, relatively to the tumor position, and range errors, concerned by the nature of the tissues the protons must traverse to reach a targeted depth. Integrating the robustness into each optimization schemes presented in this paper can be complex and will impact the quality of the results. The usual way to do it is to compute N scenarios that takes into account all types of errors that can occur before and during treatment, *e.g.*, the patient set up errors are modeled as a shift of the beamlets, giving N sets of beamlets. The next step is to determine the worst-case scenario and to minimize the corresponding objective function; repeating these two steps iteratively allows the performance to be lower-bounded. In our in-house TPS, *OpenTPS*, we have implemented this worst-case robust optimization. FISTArcPT could then be able to produce robust-optimized plans already. A robust version of our local search algorithm can be obtained by considering worst-case scenarios in a sampling-based approach. In fact, as long as the sample remains of reasonable size, our objective function can still be incrementally computed on each scenario independently. Finally, the method of mixed-integer programming will require more elaborated mathematical techniques in order to obtain robust treatment plans. Given a set of sampled scenarios, the L-shaped method [14] can be exploited to minimize expected penalties. Dealing with the inherent uncertainties of our treatment process is definitively a challenging issue, yet necessary. This is, however, out of our current scope, and left for further work.

10. Conclusions and future research directions

Arc proton therapy is an emerging cancer treatment modality, showing great potential and significant advantages over radiotherapy with x-rays and even over conventional proton therapy treatments. However, designing a valid treatment plan for this new modality defines a complex optimization problem, for which very few plan optimization methods exist in the particle therapy community. In practice,

at the time of writing, there is still no proven or approved software solution available for the clinicians and further investigations is needed to develop the treatment planning system for ArcPT.

Contributions. This paper first reviews the literature about the existing approaches to ArcPT that raise interest in the proton therapy community, namely, SPArc to which we were able to compare our results and ELO-SPAT from which we based the formulation of one of our proposed method. Next, a first formal statement of the constrained optimization problem at hand is presented. Cues in this formalization guide the search for adequate optimization algorithms. Accordingly, we propose three different approaches to solving the problem, each with a specific problem restatement and an optimization algorithm. Our new methods are tested on different realistic benchmarks and compared to the existing state-of-the-art method SPArc as well as to a conventional IMPT method. After analysis of the influence of hyper-parameters in each method, our algorithms were able to produce acceptable ArcPT treatment plans. These plans achieve good dosimetric results and delivery efficiency in general. MIP has shown its superiority compared to the other methods for small targets, as it can give guarantees on the constraints and objectives defined in the statement of a small-scale problem. However, for bigger targets, the IMPT plans and SPArc plans are still ahead which leaves some space to improve our methods depending on the direction we are taking. Anyways, each clinical case can be different and it is understandable that a method could be more appropriate than another for a specific case but improper for another case. We should stay realistic and pick the right algorithm depending on the clinical parameters at stake. Finally, we introduce the first reference benchmark to design ArcPT plans, obtained from realistic image data sets. This benchmark is made available to the research community.

Further work. The algorithms both optimize the clinical goals and minimize the delivery time in a single integrated process, thereby contrasting with SPArc. This followed path clearly suggests the use of a Multi-Criteria Optimization (MCO) that will be investigated into the future. Robustness is key in proton therapy due to its dosimetric sensitivity compared to the photon radiotherapy. Implementation of robust optimization in our ArcPT framework is therefore considered in the near future in order to make plans robust against uncertainties. Finally, this paper only considered an artificial image, *i.e.*, a patient phantom, as a proof of concept. We will apply our methods to real patients to further investigate the potential of ArcPT.

Acknowledgments

S.W. and K.S. are funded by the Walloon Region as part of the Arc Proton Therapy convention (Pôles Mecatech et Biowin). J.A.L. is a Research Associate with the F.R.S.-FNRS. Computational resources have been provided by the supercomputing facilities of the Université catholique de Louvain (CISM/UCL) 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

- [1] Amir Beck. *First-Order Methods in Optimization*. Philadelphia, PA, USA: SIAM-Society for Industrial and Applied Mathematics, 2017. ISBN: 1611974984.
- [2] Amir Beck and Marc Teboulle. “A fast iterative shrinkage-thresholding algorithm for linear inverse problems”. In: *SIAM Journal on Imaging Sciences* 2.1 (2009), pp. 183–202.
- [3] Richard H. Byrd et al. “A limited memory algorithm for bound constrained optimization”. English. In: *SIAM Journal of Scientific Computing* 16 (Sept. 1995), pp. 1190–1208. ISSN: 1064-8275. DOI: 10.1137/0916069.
- [4] Sheng Chang et al. “Feasibility study: Spot-scanning Proton Arc therapy (SPArc) for left-sided whole breast radiotherapy”. In: (June 2020). DOI: 10.21203/rs.3.rs-38925/v2.
- [5] Michaël Defferrard, Rodrigo Pena, and Nathanaël Perraudin. *PyUNLocBoX: Optimization by Proximal Splitting*. DOI: 10.5281/zenodo.1199081. URL: <https://github.com/epfl-lts2/pyunlocbox/>.
- [6] Xuanfeng Ding et al. “Have we reached proton beam therapy dosimetric limitations? – A novel robust, delivery-efficient and continuous spot-scanning proton arc (SPArc) therapy is to improve the dosimetric outcome in treating prostate cancer”. In: *Acta Oncologica* 57.3 (2018). PMID: 28774218, pp. 435–437. DOI: 10.1080/0284186X.2017.1358463.
- [7] Xuanfeng Ding et al. “Improving dosimetric outcome for hippocampus and cochlea sparing whole brain radiotherapy using spot-scanning proton arc therapy”. In: *Acta Oncologica* 58.4 (2019). PMID: 30632851, pp. 483–490. DOI: 10.1080/0284186X.2018.1555374.

- [8] Xuanfeng Ding et al. “Spot-Scanning Proton Arc (SPArc) Therapy: The First Robust and Delivery-Efficient Spot-Scanning Proton Arc Therapy”. In: *International Journal of Radiation Oncology*Biology*Physics* 96.5 (2016), pp. 1107–1116. ISSN: 0360-3016. DOI: <https://doi.org/10.1016/j.ijrobp.2016.08.049>.
- [9] Erik Engwall et al. “Fast robust optimization of proton PBS arc therapy plans using early energy layer selection and spot assignment”. In: *Physics in Medicine & Biology* 67.6 (Mar. 2022), p. 065010. DOI: [10.1088/1361-6560/ac55a6](https://doi.org/10.1088/1361-6560/ac55a6).
- [10] Wenbo Gu et al. “A novel energy layer optimization framework for spot-scanning proton arc therapy”. In: *Medical Physics* 47.5 (2020), pp. 2072–2084. DOI: <https://doi.org/10.1002/mp.14083>.
- [11] Petros Kalendralis et al. “Multicenter CT phantoms public dataset for radiomics reproducibility tests”. In: *Medical Physics* 46.3 (2019), pp. 1512–1518. DOI: <https://doi.org/10.1002/mp.13385>.
- [12] Tatsuaki Kanai et al. “Spot scanning system for proton radiotherapy”. In: *Medical Physics* 7.4 (1980), pp. 365–369. DOI: <https://doi.org/10.1118/1.594693>.
- [13] Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. “Optimization by simulated annealing”. In: *Science* 220.4598 (1983), pp. 671–680.
- [14] Gilbert Laporte and François V Louveaux. “The integer L-shaped method for stochastic integer programs with complete recourse”. In: *Operations research letters* 13.3 (1993), pp. 133–142.
- [15] Anne W. Lee et al. “International Guideline on Dose Prioritization and Acceptance Criteria in Radiation Therapy Planning for Nasopharyngeal Carcinoma”. In: *International Journal of Radiation Oncology*Biology*Physics* 105.3 (2019), pp. 567–580. ISSN: 0360-3016. DOI: <https://doi.org/10.1016/j.ijrobp.2019.06.2540>.
- [16] Xiaoqiang Li et al. “Improve dosimetric outcome in stage III non-small-cell lung cancer treatment using spot-scanning proton arc (SPArc) therapy”. In: *Radiation Oncology (London, England)* 13 (2018). DOI: <https://doi.org/10.1186/s13014-018-0981-6>.
- [17] Xiaoqiang Li et al. “The first prototype of spot-scanning proton arc treatment delivery”. In: *Radiotherapy and Oncology* 137 (2019), pp. 130–136. ISSN: 0167-8140. DOI: <https://doi.org/10.1016/j.radonc.2019.04.032>.

- [18] Gang Liu et al. “A novel energy sequence optimization algorithm for efficient spot-scanning proton arc (SPArc) treatment delivery”. In: *Acta Oncologica* 59.10 (2020). PMID: 32421375, pp. 1178–1185. DOI: 10.1080/0284186X.2020.1765415.
- [19] Gang Liu et al. “Develop An Accurate Model of Spot-Scanning Treatment Delivery Time for a Compact Superconducting Synchrocyclotron Proton System”. In: Joint AAPM & COMP Meeting. July 12, 2020. URL: <https://w3.aapm.org/meetings/2020AM/programInfo/programAbs.php?sid=8491&aid=52335>.
- [20] Gang Liu et al. “Improve the dosimetric outcome in bilateral head and neck cancer (HNC) treatment using spot-scanning proton arc (SPArc) therapy: A feasibility study”. In: *Radiation Oncology* 15 (Jan. 2020). DOI: 10.1186/s13014-020-1476-9.
- [21] Gang Liu et al. “Is proton beam therapy ready for single fraction spine SBRS? – a feasibility study to use spot-scanning proton arc (SPArc) therapy to improve the robustness and dosimetric plan quality”. In: *Acta Oncologica* 60.5 (2021). PMID: 33645429, pp. 653–657. DOI: 10.1080/0284186X.2021.1892183.
- [22] A Lomax. “Intensity modulation methods for proton radiotherapy”. In: *Physics in Medicine and Biology* 44.1 (Jan. 1999), pp. 185–205. DOI: 10.1088/0031-9155/44/1/014.
- [23] Constantin Niculescu and Lars-Erik Persson. *Convex Functions and Their Applications: A Contemporary Approach*. Jan. 2006. ISBN: 978-0-387-24300-9. DOI: 10.1007/0-387-31077-0.
- [24] David Pisinger and Stefan Ropke. “Large neighborhood search”. In: *Handbook of metaheuristics*. Springer, 2019, pp. 99–127.
- [25] Florian A. Potra and Stephen J. Wright. “Interior-point methods”. In: *Journal of Computational and Applied Mathematics* 124.1 (2000). Numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations, pp. 281–302. ISSN: 0377-0427. DOI: [https://doi.org/10.1016/S0377-0427\(00\)00433-7](https://doi.org/10.1016/S0377-0427(00)00433-7).
- [26] Mahboob ur Rehman et al. “An optimized approach for robust spot placement in proton pencil beam scanning”. In: *Physics in Medicine & Biology* 64.23 (Dec. 2019), p. 235016. DOI: 10.1088/1361-6560/ab4e78. URL: <https://doi.org/10.1088/1361-6560/ab4e78>.
- [27] G.A. Sandison et al. “Phantom assessment of lung dose from proton arc therapy”. In: *International Journal of Radiation Oncology*Biophysics*Physics* 38.4 (1997), pp. 891–897. ISSN: 0360-3016. DOI: [https://doi.org/10.1016/S0360-3016\(97\)00059-X](https://doi.org/10.1016/S0360-3016(97)00059-X).

- [28] Joao Seco et al. “Proton Arc Reduces Range Uncertainty Effects and Improves Conformality Compared With Photon Volumetric Modulated Arc Therapy in Stereotactic Body Radiation Therapy for Non-Small Cell Lung Cancer”. In: *International Journal of Radiation Oncology*Biography*Physics* 87.1 (2013), pp. 188–194. ISSN: 0360-3016. DOI: <https://doi.org/10.1016/j.ijrobp.2013.04.048>.
- [29] Edward Shaw et al. “Radiation therapy oncology group: Radiosurgery quality assurance guidelines”. In: *International Journal of Radiation Oncology*Biography*Physics* 27.5 (1993), pp. 1231–1239. ISSN: 0360-3016. DOI: [https://doi.org/10.1016/0360-3016\(93\)90548-A](https://doi.org/10.1016/0360-3016(93)90548-A).
- [30] Noah Simon et al. “A Sparse-Group Lasso”. In: *Journal of Computational and Graphical Statistics* 22.2 (2013), pp. 231–245. DOI: [10.1080/10618600.2012.681250](https://doi.org/10.1080/10618600.2012.681250).
- [31] Kevin Souris, John Aldo Lee, and Edmond Sterpin. “Fast multipurpose Monte Carlo simulation for proton therapy using multi- and many-core CPU architectures”. In: *Medical Physics* 43.4 (2016), pp. 1700–1712. DOI: <https://doi.org/10.1118/1.4943377>.
- [32] Robert R. Wilson. “Radiological Use of Fast Protons”. In: *Radiology* 47.5 (1946). PMID: 20274616, pp. 487–491. DOI: [10.1148/47.5.487](https://doi.org/10.1148/47.5.487).
- [33] Sophie Wuyckens. *ArcPT Benchmark*. Version V2. 2022. DOI: [10.14428/DVN/MZVYMU](https://doi.org/10.14428/DVN/MZVYMU).

Appendices

A. (Hyper-)parameters

In this appendix, the computing performances of each algorithm are analyzed. The hyper-parameters are identified and carefully tuned to achieve the ArcPT goals. Several metrics are therefore evaluated. First, the layer sparsity, ideally aimed at 100%, *i.e.* at most one energy layer active per beam angle. Secondly, we seek to minimize as much as possible the number of energy switch-ups as well as the dose fidelity term to follow the physician’s prescription. Finally, the homogeneity index ($HI = D5-D95$) is used to quantify the uniformity of the dose in the target. It should be as small as possible. In our problem set-up, only a single clinical objective on the target (uniform dose of 60 Gy) is implemented in the cost function. For the sake of the analysis, we show results only for the phantom data set created with 2-mm dose resolution (acceptable in clinic) and comprising 46 beams (step angle = 2°).

A.1. FISTArcPT

The FISTArcPT algorithm minimizes the cost function described in Eq. 18. This parsimonious function is built on several terms with specific purposes, each controlled by a weighting parameter. By varying the value of these parameters, different results can be obtained. The physicist can therefore fine-tune these parameters in order to achieve the primary goals of the treatment under consideration.

Analysis of weight λ . The first parameter to be analyzed is the weight of the group sparsity term λ . Recall that we aim at having no more than one energy layer active per beam, hence 100% in layer sparsity. In Fig. A1a, it can be seen, that acceptable arc sparsity is obtained from $\lambda = 1$ with good dose homogeneity. Fig. A1b shows that the dose fidelity follows the same pattern as the dose homogeneity in Fig. A1a. This behavior is expected given the fact that the dose fidelity cost is a squared norm of the difference between total dose computed and the prescription. It also shows that λ should not be too low in order to turn off non-interesting layers that can potentially reduce the number of switch-ups.

Analysis of weight ϵ . Surprisingly, ϵ has no real impact on this plan. Actually, the indicator function is not that useful in this specific case. The combination of dose fidelity, sparsity and energy sequencing already gives an acceptable arc plan. It is only when $\epsilon > 1.5$, that the plan is affected. In fact, layer sparsity decreases because the re-activation of layers starts to become more important than the removal.

Analysis of weight γ . Another important parameter to consider is the energy sequencing weight γ . Results are shown in Fig. A2 with the sparsity group weight fixed at $\lambda = 1$. This parameter was really hard to tune due to the non-convex nature of the energy sequencing term. Its impact was indeed more unpredictable. A value of $\gamma = 0.1$ was eventually picked following the results on Fig. A2. If γ is too large, the gradient descent is influenced too much by the energy sequencing term and not enough by the treatment objectives (encoded in the fidelity term) and the sparsity in layers is difficult to obtain. While if γ is too small, the energy sequencing term will not have any impact anymore and the energy layer switching time will not be optimized.

From those results, we decided to fix values for the weighting parameters of the objective function for the next experiences. Table A1 summarized the chosen parameter values.

A.2. Local Search

The local search algorithm, even though differing radically from the FISTArcPT, has also several inputs to tweak and orientate the search in the space of solutions. Due to the non-deterministic character of LS, each experiment was run 5 times. Average and standard deviations over these 5 runs are therefore reported.

Analysis of weight γ . As FISTArcPT, the local search algorithm has a weight controlling the importance of the energy sequencing term in the objective function. Fig. A3 illustrates the impact of γ weight on the number of switch-ups and fidelity term. A value $\gamma = 50$ seems to be the best compromise to have an acceptable ELST and homogeneous dose in the target.

Minimum number of active energy layers per beam. We can set a minimum number of energy layers to be activated in each beam as a hard constraint in the configuration of the LS algorithm. Table A2 shows a comparison where the minimum number of active layers in a beam is zero and situation like FISTArcPT, where the minimum is set to one. It is obvious that constraining the algorithm to activate at least one energy layer per control point is really restrictive and suppresses degrees of freedom resulting in higher cost, worse HI, and worse ELST.

Maximum number of active energy layer per beam. As we have set a minimum of active energy layers per beam, we can also set a superior limit on this number. Table A3 summarizes the results. Allowing more than one layer active per beam gives better dosimetric outcomes and results in a decreased fidelity cost but at the expense of an increasing energy layer switching time.

Initial temperature and cooling factor. This couple of parameters influence the behavior of local search and are used to play on the diversification/intensification trade-off. The higher (*resp.* lower) the initial temperature (*resp.* cooling factor), the more iterations between two resets and, therefore, the higher diversification. On the opposite, a smaller initial temperature and a higher cooling factor would increase the intensification. The initial temperature and cooling factor were set according to previous knowledge to 5 and 0.7 respectively.

From the results obtained in this section, we decided to fix the parameters for the next experiences. Table A4 reports their values.

B. Convergence

For FISTArcPT, the stopping criterion used in the optimization is the maximum number of iterations. We set it to 1000 to get a good quality and deliverable plan and achieve good layer sparsity. Fig. B1 displays the result. Except the indicator function, all terms involved in the objective function described in Eq. 18 are plotted. The total cost in general shows a rapid decrease and then a smooth convergence and is mainly influenced by the group sparsity term. The energy sequencing term is more chaotic but still converges after a few hundreds of iterations.

Regarding the local search algorithm, the objective function to be minimized (Eq. 22) is simply a quadratic function of the difference between the total dose delivered to the target and the prescription (omitting the OARs) to which the ELST and sparsity objectives are added. Fig. B2 illustrates the evolution of the cost function for 5 runs. The stopping criterion is a time limit set to 5 minutes. As expected, the local search is really chaotic at the beginning, exploring solutions in a large neighborhood but little by little, converges smoothly towards an optimal solution along with the decrease of the temperature of the system. Compared to FISTA, the convergence is certainly slower but a rapid solution with specific properties could potentially be used as input to FISTA.

C. Problem scaling

The scale of the problem constitutes a major issue in ArcPT optimization. Depending on the resolution and size of the beamlet matrix, the algorithms will behave in different ways. We decided to pick 9 data samples and to analyze them: the two extremes and the middle one, *i.e.* phantoms with 1, 3, and 5 mm resolution and with a plan composed of 10, 31, and 91 beams. For the sake of the comparison, we fixed the maximum number of iterations of FISTArcPT to 1000. This way, each

test run of FISTA determines the time limit to be set in the LS and MIP algorithm input parameters (though MIP could stop before the time limit if the optimal solution was found).

As it could be expected, each algorithm behaves in a very different way and this is summarized in Tables C1-C2-C3. FISTArcPT gets better results in terms of dose homogeneity (D5 - D95) for lower resolution data samples and smaller number of beams while on the other side, LS-ArcPT behaves in a complete opposite way with better results at high resolution with a high number of beams. For both algorithms, the ELST increases with the number of beams in a logical way and are quite compatible with each other. The MIP algorithm, on the other hand, does not follow the same pattern. Indeed, the maximum number of energy upwards switches has been constrained to 0 instead of being left as an optimization objective, hence the very low ELST for each test set. As it can be observed in Table C3, the MIP gives excellent results in terms of dose homogeneity within the target volume and this, for any of the test sets. The 5-mm resolution data sets could achieve a better homogeneity index compared to other resolution sets because the MIP let us push the constraints on the target coverage harder ($[p_{\min T} = 60, p_{\max T} = 60.5]$ Gy instead of $[58,62]$ Gy for the others). It can also be noticed that the boxes regarding the 1-mm resolution data set with 10 beams in the arc are empty. The reason is that the MIP could not find any feasible solution with the constraints given even though many combinations were tried.

In terms of optimization time, FISTArcPT behaves as expected because, a higher number of beams translates into a higher number of beamlets and then a more computationally expensive gradient. The MIP, for its part, was able to give an optimal solution in a shorter time than the duration limit set by the FISTA algorithm and this, for any of the test sets. Obviously, the resolution also plays a role in optimization time with higher resolution, *i.e.* higher number of voxels that translates in larger optimization time. Despite the fact that this claim is valid for any of the presented optimization algorithm, the MIP seems to be the one that could be the most impacted with the time exponentially with the problem size.

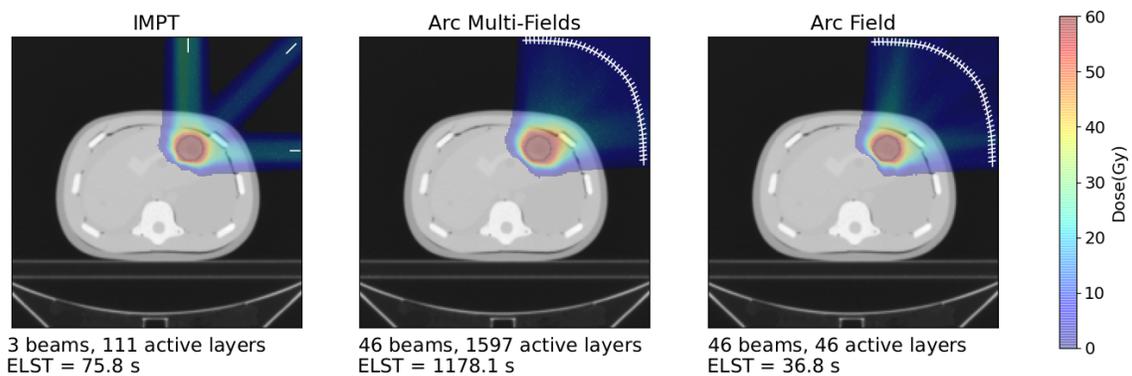


Figure 1: PT modality and delivery comparison



Figure 2: Proton therapy Room at Beaumont Hospital with IBA arc prototype

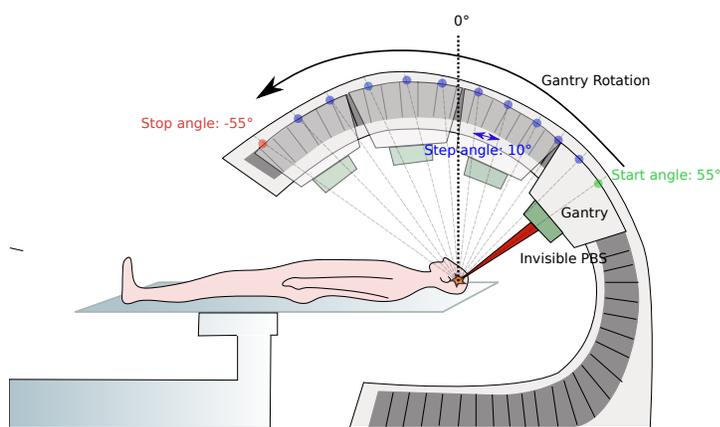


Figure 3: Arc proton therapy principle

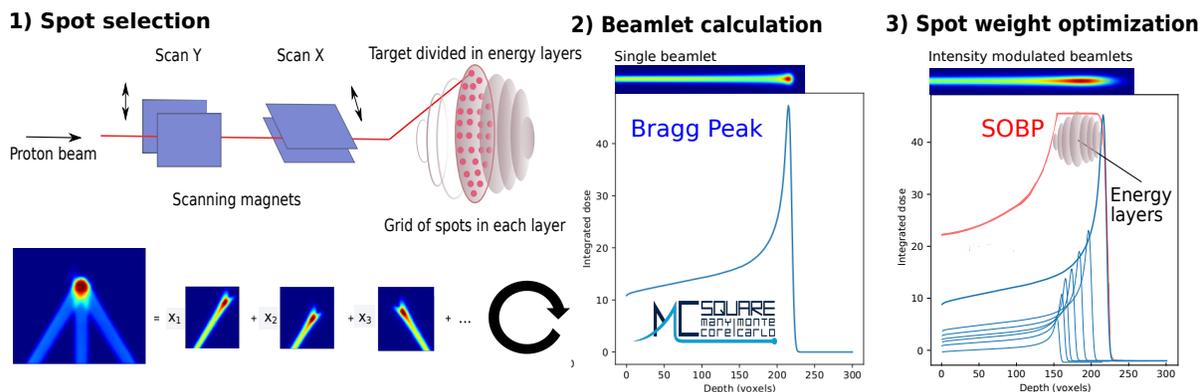


Figure 4: Typical PT treatment plan generation steps.

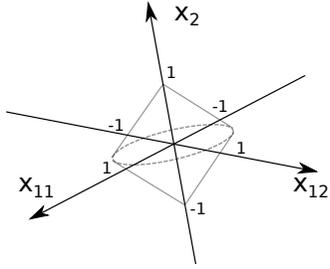


Figure 5: Group Lasso penalty: $\|x_1\| + |x_2| = 1$

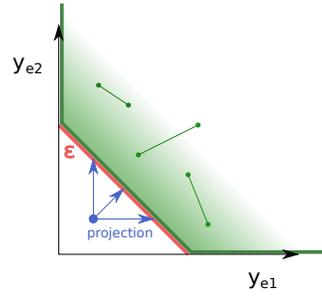


Figure 6: Simplified representation of set C and its convexity. Any pair of points (green) $\in C$, can be connected by a line that lies in C .

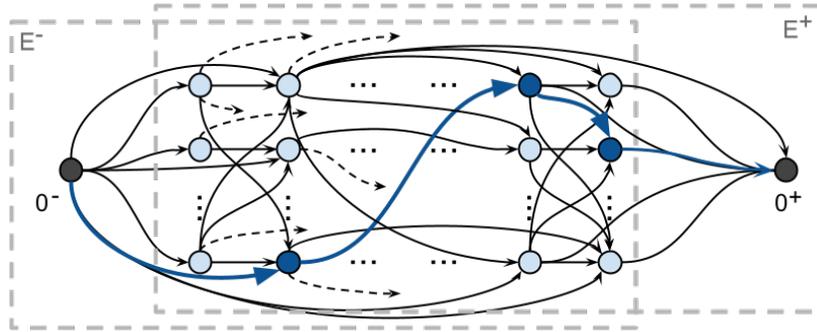


Figure 7: The directed graph resulting from the E^{edges} edges construction. Source and sink nodes are represented as the left and right circles. A path in the digraph is necessarily a valid sequence of activated energy layers. In this simplified example, bold edges represent a possible unique path (from source node 0^- to sink node 0^+). in which the dark blue nodes are the activated ELs, and where the first beam has no activated EL.

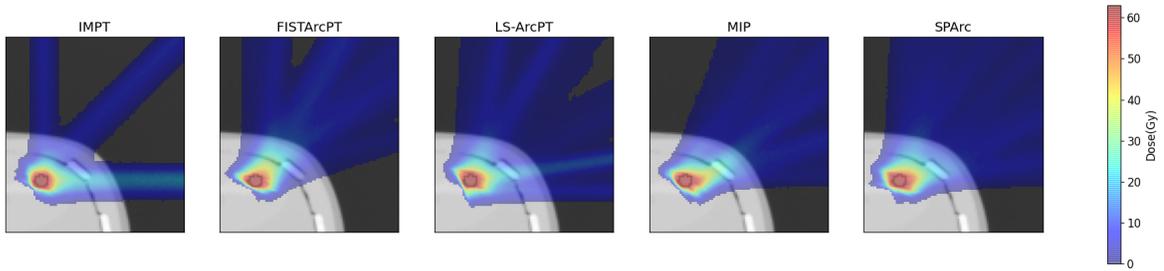


Figure 8: (Color for print) Optimized dose distributions (2mm resolution) obtained with the different modalities/algorithms - Small target

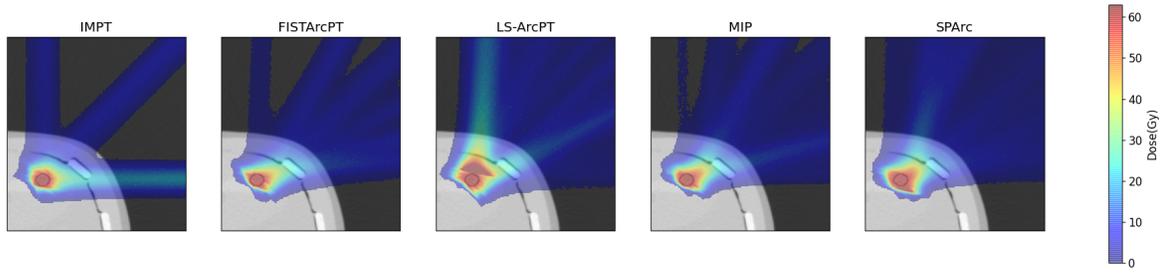


Figure 9: (Color for print) Optimized dose distributions (1mm resolution) obtained with the different modalities/algorithms - Small target

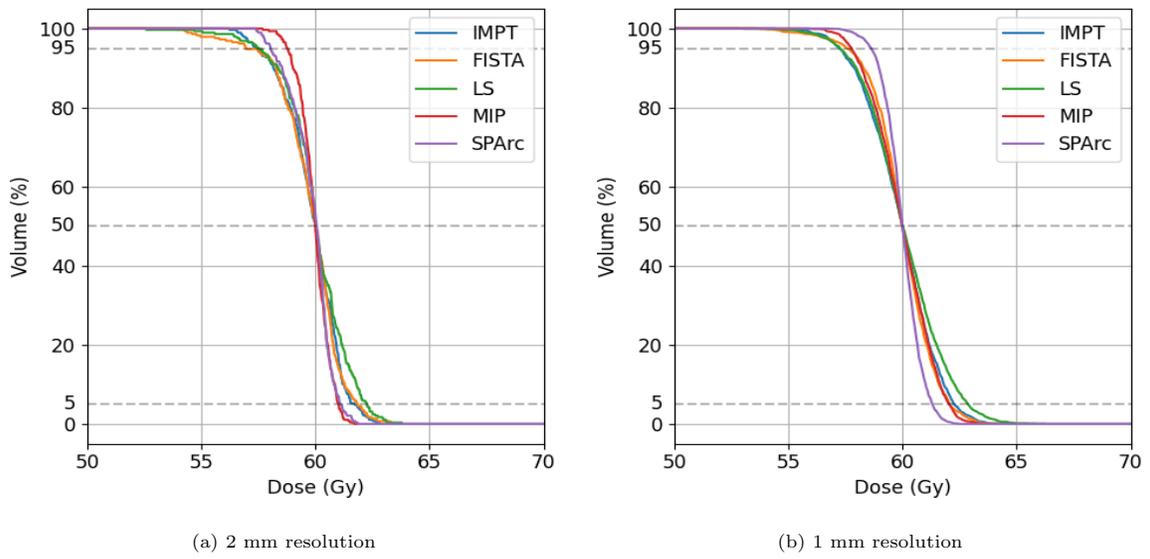


Figure 10: (Color for print) DVH comparison between the different algorithms used for the small target

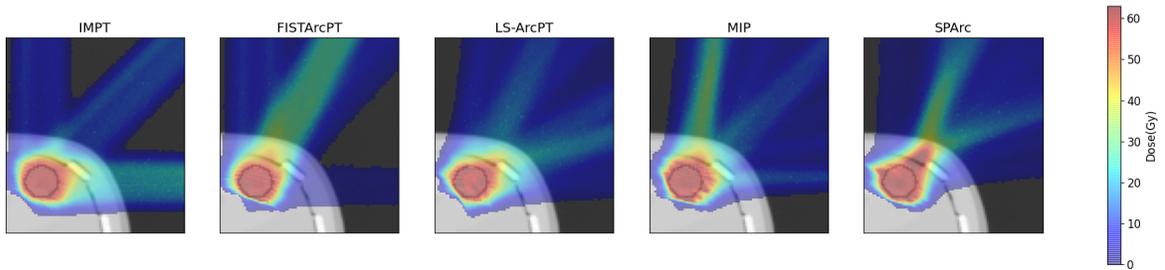


Figure 11: (Color for print) Optimized dose distributions (2mm resolution) obtained with the different modalities/algorithms - Expanded target

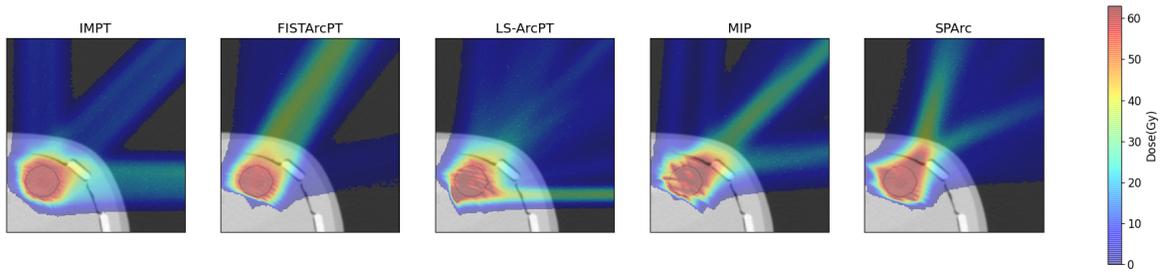


Figure 12: (Color for print) Optimized dose distributions (1mm resolution) obtained with the different modalities/algorithms - Expanded target

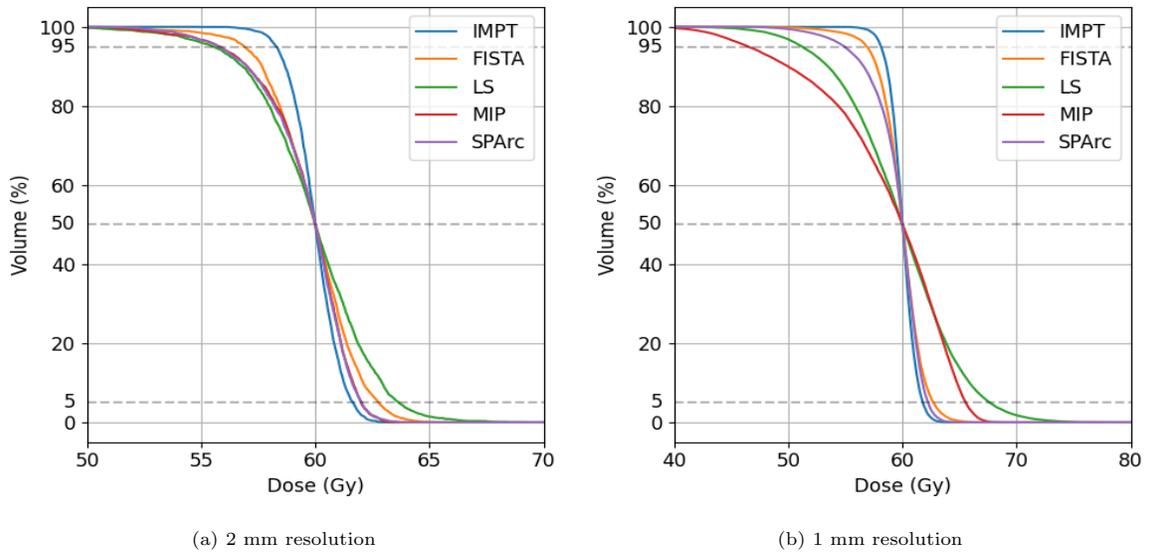


Figure 13: (Color for print) DVH comparison between the different algorithms used for the expanded target

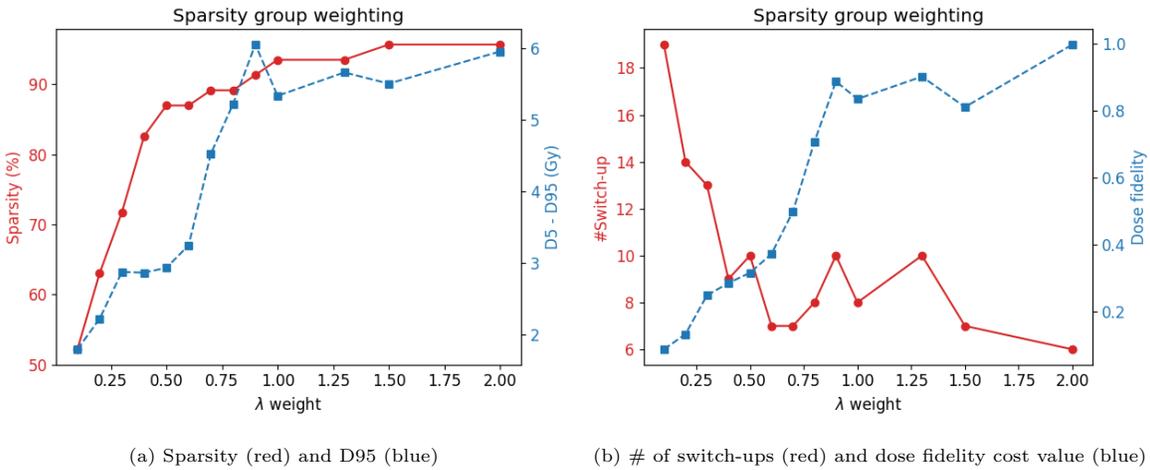
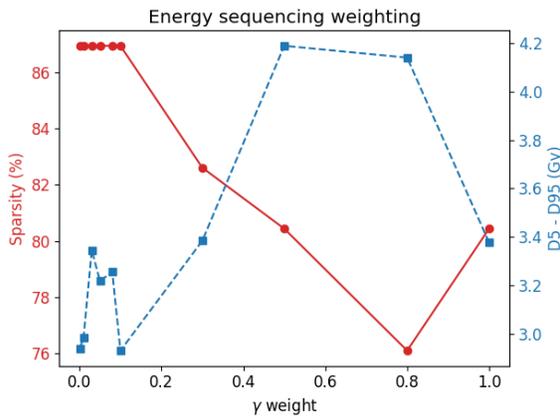
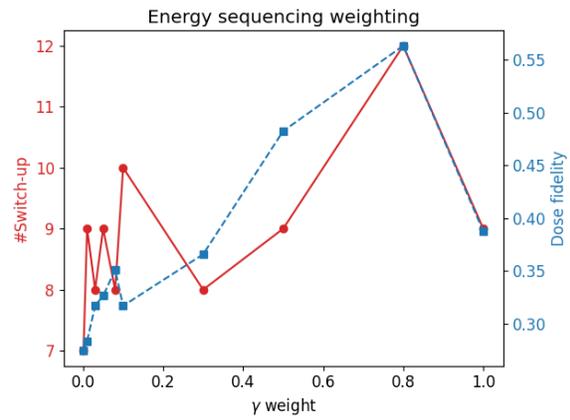


Figure A1: (Color for print) FISTArcPT metric versus sparsity group weighting λ

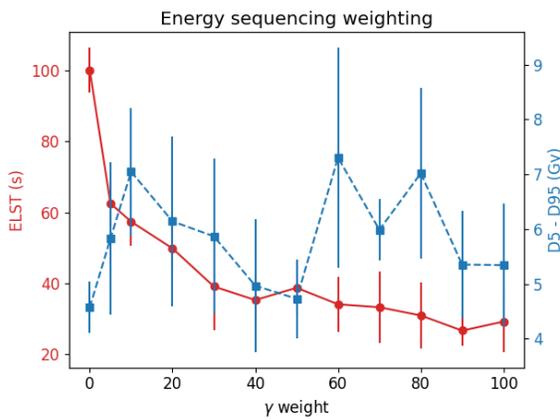


(a) Sparsity (red) and HI (blue)

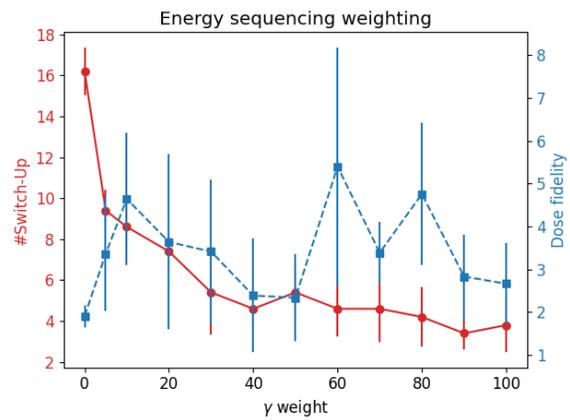


(b) # of switch-ups (red) and fidelity cost (blue)

Figure A2: (Color for print) FISTArcPT metric versus energy sequencing weighting γ



(a) ELST (red) and HI (blue)



(b) # of switch-ups (red) and fidelity cost (blue)

Figure A3: (Color for print) LS metric versus energy sequencing weighting γ for 5 runs

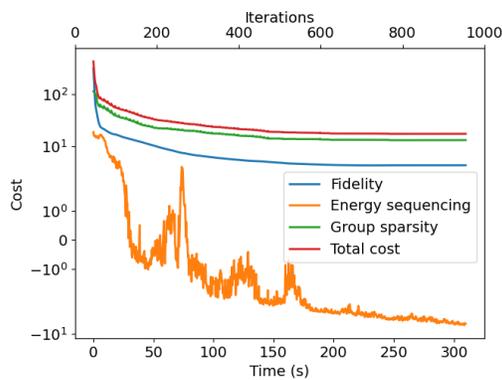


Figure B1: (Color for print) Convergence of FISTArcPT objective function

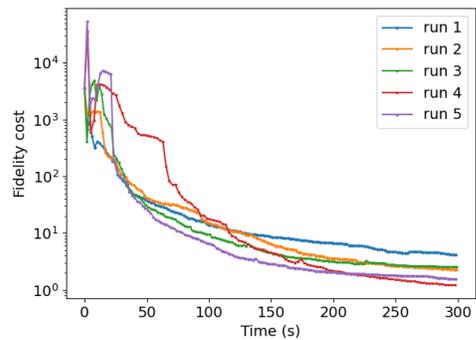


Figure B2: Convergence of 5 runs of LS-ArcPT with the same input parameters

CT resolution		1 mm	2 mm	3 mm	4 mm	5 mm
Number of voxels		20,010,375	2,469,888	720,896	304,128	150,176
Step angle	Number of spots	Beamlet computation time and memory				
1°: 91 beams	13,105	5373.3 s 10.5 GB	600.1 s 1.3 GB	185.4 s 395 MB	78.9 s 169 MB	41.5 s 88 MB
2°: 46 beams	6,635	2429.8 s 5.3 GB	297.0 s 668 MB	91.3 s 200 MB	38.5 s 86 MB	20.8 s 45 MB
3° : 31 beams	4,469	1770.8 s 3.6 GB	204.6 s 450 MB	62.0 s 134 MB	27.3 s 58 MB	13.3 s 30 MB
5°: 19 beams	2,622	992.2 s 2.1 GB	117.5 268 MB	36.8 s 80 MB	16.7 s 34 MB	8.4 s 18 MB
10° : 10 beams	1,453	549.8 s 1.2 GB	64.5 s 147 MB	20.1 s 44 MB	8.7 s 19 MB	4.3 s 10 MB

Table 1: Benchmarks characterization of beamlets computation time and file size associated with different CT resolution and step angles

		IMPT	ArcPT
Small target	#ELs	39	361
	#spots	1805	6635
Expanded target	#ELs	25	558
	#spots	400	29269

Table 2: Degrees of freedom available to optimizer

Metric/Method	IMPT		FISTA		LS		MIP		SPArc	
Resolution (mm)	2	1	2	1	2	1	2	1	2	1
Planned dose quality										
$ \Delta(D95, D5) $ (Gy)	4.2	5.0	4.0	4.6	4.8	5.5	2.2	4.0	2.9	2.7
D98 (Gy)	56.8	56.7	55.0	56.3	56.5	56.5	58.5	57.3	57.7	58.0
CI_{RTOG}	1.4	1.1	1.2	1.4	1.3	3.7	1.1	1.2	1.3	1.9
Delivery efficiency										
#switch-ups	2	2	4	5	2	2	0	0	2	2
#switch-down	22	22	10	9	14	9	8	6	13	12
ELST (s)	24.2	24.2	28	32.9	19.4	16.4	4.8	3.6	18.8	18.2
Sparsity (%)	0	0	95.6	93.5	100	100	100	100	100	100
Degrees of freedom										
#ELs activated	25	25	48	48	42	42	40	39	46	46
Non zero spots	400	400	1198	1214	131	228	171	217	1163	1225
Computational performance										
Optimization time (s)	2	12	2000	14000	300	3600	3600	3600	714	4423

Table 3: Metrics comparison between the different methods used for small target

Metric/Method	IMPT		FISTA		LS		MIP		SPArc	
Resolution (mm)	2	1	2	1	2	1	2	1	2	1
Planned dose quality										
$ \Delta(D95,D5) $ (Gy)	3.3	3.7	6.0	5.8	8.4	18.4	6.4	23.5	6.5	7.5
D98 (Gy)	57.7	57.5	55.5	55.3	53.8	48.8	53.8	43.6	54.1	52.6
CI_{RTOG}	1.3	1.4	1.1	1.1	1.2	1.5	1.3	2.3	1.7	1.7
Delivery efficiency										
#switch-ups	2	2	11	14	2	2	3	2	2	2
#switch-down	35	36	25	24	22	18	18	7	20	20
ELST (s)	32.0	32.6	75.5	91.4	24.2	21.8	27.3	15.2	23.0	23.0
Sparsity (%)	0	0	84.8	80.4	100	100	100	100	100	100
Degrees of freedom										
#ELs activated	38	39	60	60	46	46	29	12	46	46
Non zero spots	1897	1805	2949	3139	457	511	433	272	2539	2474
Computational performance										
Optimization time (s)	19	181	20000	76000	3600	3600	12000	12000	2273	14423

Table 4: Metrics comparison between the different methods used for expanded target

Parameter names	λ	γ	ϵ
Value	1.	0.1	0.1

Table A1: Tuned parameters in FISTArcPT algorithm

	Minimum 0 active layer in each beam	Minimum 1 active layer in each beam
Fidelity cost	1.29 ± 0.27	1.66 ± 0.58
D5 - D95 (Gy)	7.04 ± 1.21	7.79 ± 1.66
ELST (s)	57.38 ± 6.86	96.66 ± 10.66
#Switch-Up	8.60 ± 1.20	15.00 ± 2.10

Table A2: Impact of minimum number of active layer per beam for 5 runs

	Max 1 active layer in each beam	Max 2 active layers in each beam	Max 3 active layers in each beam
Fidelity cost	1.29 ± 0.27	1.19 ± 0.57	0.69 ± 0.33
D5 - D95 (Gy)	7.04 ± 1.21	6.62 ± 1.55	4.81 ± 1.30
ELST (s)	57.38 ± 6.86	80.56 ± 13.36	95.36 ± 19.64
#Switch-Up	8.60 ± 1.20	11.20 ± 2.04	12.80 ± 2.79
Sparsity (%)	100.00 ± 0.00	49.13 ± 7.48	59.57 ± 10.69

Table A3: Impact of maximum number of active layer per beam for 5 runs

Parameter names	γ	minEL in beams	maxEL in beams	Target weight	Max spot intensity
Value	50.0	0	1	100.0	20

Table A4: Tuned parameters in LS algorithm

	Resolution	# beams		
		10	31	91
ELST (total)	1 mm	22.5''	32.9''	75.6''
	3 mm	29.8''	34.7''	40.2''
	5 mm	21.9''	34.1''	40.2''
Optimization time	1 mm	2435''	2595''	3640''
	3 mm	180''	500''	1800''
	5 mm	60''	160''	860''
D5 - D95 (Gy)	1 mm	5.0	4.7	5.5
	3 mm	3.1	3.4	4.0
	5 mm	1.6	2.0	2.6
Layer	1 mm	60.0	90.32	92.31
	3 mm	50.0	83.87	98.9
Sparsity (%)	5 mm	60.0	90.32	96.7

Table C1: Summary of FISTArcPT experiences

	Resolution	# beams		
		10	31	91
ELST (total)	1 mm	25.0''	41.4''	56.6''
	3 mm	14.6''	16.4''	82.8''
	5 mm	13.4''	32.3''	55.4''
D5 - D95 (Gy)	1 mm	9.5	4.6	5.4
	3 mm	14.2	6.3	2.9
	5 mm	7.8	1.8	1.4

Table C2: Summary of LS experiences

	Resolution	# beams		
		10	31	91
ELST (total)	1 mm	/	3"	169.5"
	3 mm	2.4"	16.4"	3.6"
	5 mm	8.5"	3.6"	3.6"
Optimization time	1 mm	/	2041.1"	1357.8"
	3 mm	22.3"	65.8"	398.6"
	5 mm	2.9"	12.5"	97.8"
D5 - D95 (Gy)	1 mm	/	4.1	3.8
	3 mm	5.1	5.6	6.2
	5 mm	0.7	1.2	2.1
Layer	1 mm	/	100.0	78.0
	3 mm	100.0	87.1	100.0
Sparsity (%)	5 mm	80.0	100	100

Table C3: Summary of MIP experiences