



A stochastic approach to full inverse treatment planning for charged-particle therapy

Marc C. Robini¹ · Feng Yang² · Yuemin Zhu¹

Received: 31 October 2018 / Accepted: 7 March 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

Charged-particle therapy is a rapidly growing precision radiotherapy technique that treats tumors with ion beams. Because ion-beam delivery systems have multiple degrees of freedom (including the beam trajectories, energies and fluences), it can be extremely difficult to find a treatment plan that accurately matches the dose prescribed to the tumor while sparing nearby healthy structures. This inverse problem is called *inverse treatment planning* (ITP). Many ITP approaches have been proposed for the simpler case of X-ray therapy, but the work dedicated to charged-particle therapy is usually limited to optimizing the beam fluences given the trajectories and energies. To fill this gap, we consider the problem of simultaneously optimizing the beam trajectories, energies, and fluences, which we call *full ITP*. The solutions are the global minima of an objective function defined on a very large search space and having deep local basins of attraction; because of this difficulty, full ITP has not been studied (except in preliminary work of ours). We provide a proof of concept for full ITP by showing that it can be solved efficiently using simulated annealing (SA). The core of our work is the incremental design of a state exploration mechanism that substantially speeds up SA without altering its global convergence properties. We also propose an original approach to tuning the cooling schedule, a task critical to the performance of SA. Experiments with different irradiation configurations and increasingly sophisticated SA algorithms demonstrate the benefits and potential of the proposed methodology, opening new horizons to charged-particle therapy.

Keywords Simulated annealing · Markov chain Monte Carlo (MCMC) · Inverse treatment planning · Charged-particle therapy · Hadrontherapy

✉ Marc C. Robini
marc.robini@creatis.insa-lyon.fr

✉ Feng Yang
yalltroy@gmail.com

¹ CREATIS Laboratory (CNRS Research Unit UMR5220 and INSERM Research Unit U1206), INSA Lyon, 69621 Villeurbanne, France

² National Library of Medicine, National Institute of Health, Bethesda, MD 20894, USA

1 Introduction

1.1 Motivations and background

The goal of external beam radiotherapy is to irradiate tumors without damaging surrounding healthy tissues and nearby organs at risk (OARs). In particular, charged-particle therapy, also known as hadrontherapy, is a rapidly growing technique that uses pencil beams of protons or heavier ions (such as carbon) to deliver a prescribed dose to the target tumor [1–5]. The advantage of ion beams over X-rays is their good selectivity: they present low or negligible lateral scattering, and their depth-dose profile has a sharp maximum—the *Bragg peak*—whose position can be adjusted by changing the kinetic energy of the incident ions.

The first step in a radiotherapy treatment is to delineate the target tumor on CT or MRI scan images, with safety margins accounting for setup uncertainties. This defines the so-called *planning target volume* (PTV). The OARs are defined in the same way and their union is referred to as the OAR region. Then, based on the tumor characteristics and the OARs, a radiation oncologist prescribes the dose to be delivered to the PTV and chooses the fractionation scheme, that is, the number of sessions (or *fractions*) of the treatment and their frequency. For example, in the case of carbon, the prescribed dose is typically delivered in 4 to 20 fractions over 1 to 6 weeks.

We focus on the next step, namely *inverse treatment planning* (ITP), which is to find a treatment plan that delivers the prescribed dose to the PTV while sparing the OAR region. Because charged-particle therapy has many degrees of freedom (including the beam trajectories, energies and fluences), ITP is considered intractable in its full generality and is usually partially solved by optimizing the beam fluences given the trajectories and energies. This suboptimal approach is called spot-intensity map optimization (SIO) and is similar to fluence map optimization (FMO) in X-ray therapy. The dominant approaches to SIO and FMO are linear programming [6–10] and quadratic programming [11–21]. Other approaches include convex programming [22,23], mixed integer programming [24], minimax stochastic programming [25], simulated annealing [26], and genetic algorithms [27–29].

There are currently three more general alternatives to SIO in the charged-particle therapy literature: joint optimization of the beam energies and fluences [30], beam orientation optimization (BOO) with decoupled SIO [31,32], and joint beam orientation and spot-intensity map optimization [33]. In [30], the authors propose an iterative mixed integer programming algorithm to simultaneously optimize the beam energies and fluences while limiting the number of distinct energies used for each beam orientation. In [31], the beam orientations are optimized by local neighborhood search, with SIO performed for each visited configuration using linear programming. This latter approach is improved in [32], where different global optimization algorithms are used to provide a good starting point for local neighborhood search. Finally, in [33], the beam orientations and the spot intensities are optimized simultaneously using a fast iterative shrinkage-thresholding algorithm to minimize an objective function consisting of three terms: a dose-fidelity term measuring the deviation to the prescribed dose, an ℓ_1 -sparsity penalty limiting the number of active spots, and a group-sparsity penalty controlling the number of active beam orientations.

Note that the BOO problem has been widely addressed for X-ray therapy (see, e.g., [34] for a comprehensive introduction). In this case, the general ITP problem is to select the beam orientations (usually 5–9 coplanar beam angles) and to compute their fluence maps. Again, two BOO strategies can be identified. The first one is to perform FMO independently for each visited beam orientation configuration; that is, ITP is formulated as the minimization of an expensive objective function whose evaluation requires solving an FMO problem.

The advantage is that virtually any general-purpose optimization algorithm can be used to search for the optimal beam orientations: multistart gradient-based local search [35], response surface [36], nested partitions [37], local neighborhood search [38], greedy set cover [39], pattern search [40], simulated annealing [41–44], genetic algorithms [45–47], and particle swarm [48] (see also [49,50] for a comparison of simulated annealing, genetic, and cross-entropy methods). The second, less common strategy merges BOO and FMO by simultaneously optimizing the beam orientations and fluences; the ITP objective is then much cheaper to evaluate, but the search space is much larger. Examples in this category include mixed integer programming [51–53], genetic algorithms [54,55], and convex optimization with a sparsity-inducing penalty [56].

Coming back to charged-particle therapy, we face the additional problem of optimizing the Bragg peak positions via the beam energies. Moreover, for greater generality, we want a framework in which the possible beam trajectories are not necessarily constrained to preferential orientations. To our knowledge, apart from preliminary work of ours [57], the joint optimization of the beam trajectories, energies and fluences has not yet been considered.

1.2 Overview

We propose an efficient and theoretically sound simulated annealing (SA) approach to what we call the *full* ITP problem, which is to find the optimal trajectories, energies and fluences of a given number K of ion beams with fixed transverse profile.

The space of all possible treatment plans is $(\mathcal{T} \times \mathcal{E} \times \mathcal{N})^K$, where \mathcal{T} is the set of allowed beam trajectories (the number of which is much larger than K) and \mathcal{E} and \mathcal{N} are the sets of possible energies and fluences. These three sets are constrained by the ion-beam delivery system; in addition, \mathcal{T} may be adapted to the tumor characteristics using expert knowledge. Note that K can be either the number of beams of the entire treatment plan (meaning ITP is performed prior to fractionation) or the number of beams in a single-fraction plan (meaning the fractions are optimized independently of each other, possibly varying K and the prescribed dose from one fraction to the other). The latter option is preferable not only for computational complexity reasons, but also to account for interfraction variations of the patient geometry (e.g., tumor growth or shrinkage, tumor and organ motion, weight loss or gain). For simplicity, the term *treatment plan*, when used alone, will refer to the entire plan or a single fraction without distinction.

We define an optimal treatment plan to be a global minimum of an objective function $U(\omega)$ measuring the conformity between the prescribed dose and the dose delivered by the input treatment plan ω . Because the search space is very large (much larger than in the partial ITP problems discussed in the previous section), and because of the intricate relationship between the dose delivered and the beam trajectories and energies, the objective has deep local basins,¹ which makes its minimization challenging and excludes greedy approaches.

The behavior of an SA algorithm is governed by a communication mechanism and a cooling schedule. The communication mechanism is a Markov matrix whose entries are the probabilities of the moves in the search space, and the cooling schedule is a decreasing sequence of temperatures controlling the acceptance rate of uphill moves (i.e., moves that increase the objective). The conditions for global convergence are quite mild: the communi-

¹ The local minima and their basins depend on how the search space is explored. By claiming that the objective has deep local basins, we actually mean that a landscape defined by the objective and a tractable exploration mechanism will likely have many deep (possibly poor) local minima. This fact will be confirmed by our experiments.

cation mechanism must be irreducible and have symmetric support, and the cooling schedule must be logarithmic or exponential depending on whether we are interested in asymptotic or finite-time convergence [58,59]. Therefore the key points in designing an efficient SA algorithm are, first, to construct a communication mechanism striking a good balance between rate of convergence and computational cost and, second, to properly tune the cooling schedule. The core of our work is the construction of increasingly sophisticated communication mechanisms for full ITP. The final strategy to explore the search space is based on temperature-dependent moves that select the trajectories uniformly at random (u.a.r.) and find optimal energy-fluence pairs in randomly chosen partition blocks of $\mathcal{E} \times \mathcal{N}$. We also devise a new method for selecting the initial and final temperatures of the cooling schedule: we build a specific auxiliary SA chain from which we can accurately estimate the temperatures achieving given acceptance rates at the beginning and end of the optimization process.

Numerical experiments with different irradiation configurations show that the proposed exploration strategy yields a computational speed-up of several orders of magnitudes over simply designed communication. We also show that our full ITP approach significantly outperforms SIO in terms of treatment plan quality and offers the possibility of substantially reducing the number of spots needed for tumor coverage. However, the assessment of potential clinical benefits is beyond the scope of this paper; our purpose here is to thoroughly introduce a new strategy for full ITP in charged-particle therapy.

This paper is organized as follows. We formulate the full ITP problem in Sect. 2 and review the relevant SA theory in Sect. 3. Our exploration strategy is developed in Sects. 4 and 5 (Sect. 5 focuses on implementation and computational complexity and can be skipped on first reading). The issue of tuning the cooling schedule is addressed in Sect. 6, followed by numerical experiments in Sect. 7. Concluding remarks are given in Sect. 8.

1.3 Notation

The abbreviations and main notation used in this paper are summarized in Table 1. Generally, we use bold lower-case roman letters (e.g., \mathbf{v}) for vectors, and caligraphic upper-case letters (e.g., \mathcal{A}) for sets. The cardinality (or size) of a set \mathcal{A} is denoted by $|\mathcal{A}|$. For any $a, b \in \mathbb{Z}$ such that $a \leq b$, we let $[a \dots b]$ denote the set of integers in the interval $[a, b]$.

2 The full inverse treatment planning problem

2.1 State space

An ion beam is represented by a trajectory (that is, a ray defining its center axis), an energy $e \in \mathbb{R}_+^*$, and a number of particles, or fluence, $n \in \mathbb{N}$. A *treatment plan* is a finite set of beams $(\mathbf{b}_1, \dots, \mathbf{b}_K) \in (\mathcal{T} \times \mathbb{R}_+^* \times \mathbb{N})^K$, where \mathcal{T} is an index set labelling the allowed beam trajectories (for convenience, we make no distinction between an index $\tau \in \mathcal{T}$ and the trajectory indexed by τ); we call \mathcal{T} the *irradiation geometry*. In practice, because of the technical limitations of ion-beam delivery systems, \mathcal{T} is finite (although much larger than K), and so are the possible values for the energy and fluence.

Since a beam is only useful if its Bragg peak lies in the PTV, we assume that every trajectory intersects the PTV, and we define the set of possible energies for a trajectory τ to be

$$\mathcal{E}(\tau) := \{l\delta_e\}_{l \in \mathbb{N}} \cap [e_{\min}(\tau) - \delta_e, e_{\max}(\tau) + \delta_e], \quad (2.1)$$

Table 1 Abbreviations and main notation

BOO	Beam orientation optimization
FMO	Fluence map optimization
FWHM	Full width at half maximum
ITP	Inverse treatment planning
OAR	Organ at risk
PTV	Planning target volume
SA	Simulated annealing
SIO	Spot-intensity map optimization
u.a.r.	Uniformly at random
MeV u ⁻¹	Energy per nucleon
Gy	Unit of absorbed dose (Gray), 1 Gy = 1 J kg ⁻¹
\mathbb{R}_+^*	Set of positive reals
\mathbb{N}^*	Set of positive integers
$\lfloor \cdot \rfloor, \lceil \cdot \rceil, \lceil \cdot \rceil$	Floor, ceiling and nearest integer functions
K	Number of treatment beams
\mathcal{T}	Irradiation geometry (i.e., set of allowed beam trajectories)
$\mathcal{E}(\tau)$	Set of possible energies for a beam with trajectory τ , (2.1)
\mathcal{N}	Set of possible beam fluences, (2.2)
δ_e, δ_n	Energy and fluence increments
n_{\max}	Maximum fluence
$\mathbf{b} := (\tau, e, n)$	Ion beam with trajectory τ , energy e , and fluence n
$\omega := (\mathbf{b}_1, \dots, \mathbf{b}_K)$	Treatment plan (i.e., set of K ion beams)
$\mathcal{T} \otimes \mathcal{E}$	Set of possible trajectory-energy pairs, (2.3)
Ω	Set of feasible treatment plans, or state space, (2.3)
$\overline{ \mathcal{E} }$	Average number of energies per trajectory, (2.7)
\mathcal{V}	Set of voxels representing the PTV-OAR region
U	Objective function, (2.9)
D_\star	Dose plan
$D(v, \mathbf{b})$	Physical dose deposited in a voxel v by a beam \mathbf{b} , (5.2)
$\{\mathbf{b}(i)\}_i$	Ray discretization of \mathbf{b}
$ \mathcal{I} $	Number of rays modelling \mathbf{b}
\mathcal{D}	Set of fluence-normalized doses, (5.4)
R	Beam-to-voxel cross-section ratio, (5.23)
Q_T	SA transition matrix, (3.7)
M	Horizon (i.e., length) of the cooling schedule
$(T_m, M)_m$	Finite cooling schedule with horizon M
T_{\max}, T_{\min}	Initial and final temperatures
$\alpha_{m, M}$	Exponential cooling exponent, (3.13)
ξ	Number of constant-temperature stages
Θ_1	Communication with uniform candidate beams (Sect. 4.1)
$\omega_{\setminus k}(\tau, e, n)$	Treatment plan ω with the k th beam \mathbf{b}_k replaced by (τ, e, n)
$n_{\omega_{\setminus k}}(\tau, e)$	Optimal fluence of the k th beam, (4.3)
$\Theta_{2, \eta}$	Communication with optimal fluence moves (Sect. 4.2)

Table 1 continued

$\{\mathcal{E}_{p,j}(\tau)\}_j$	Partition of the energies into p interlaced subsets, (4.9)
$\Theta_{3,p,\eta}$	Communication with semi-optimal energy-fluence moves (Sect. 4.3)
$\mathcal{V}(\mathbf{b})$	Set of voxels in \mathcal{V} intersected by a beam \mathbf{b}
N_d	Average number of floating point operations for computing the entry and exit depths of a ray into a voxel

where δ_e is a positive real constant (independent of τ) and $e_{\min}(\tau)$ and $e_{\max}(\tau)$ are the energies that position the Bragg peak along τ at the proximal and distal edges of the PTV, respectively. The set of possible values for the fluence does not depend on the trajectory and is given by

$$\mathcal{N} := \{l\delta_n\}_{l \in \mathbb{N}} \cap [0 \dots n_{\max}], \quad (2.2)$$

where δ_n and $n_{\max} \gg \delta_n$ are positive integer constants.

A treatment plan $(\mathbf{b}_1, \dots, \mathbf{b}_K)$ is said to be *feasible* if $\mathbf{b}_k := (\tau_k, e_k, n_k) \in \mathcal{T} \times \mathcal{E}(\tau_k) \times \mathcal{N}$ for every k . The set of feasible treatment plans, or *state space*, is

$$\Omega := (\mathcal{T} \otimes \mathcal{E}) \times \mathcal{N}^K, \quad \mathcal{T} \otimes \mathcal{E} := \bigcup_{\tau \in \mathcal{T}} \{\tau\} \times \mathcal{E}(\tau). \quad (2.3)$$

It is important to note that Ω is considerably larger than the state spaces associated with SIO and BOO. In the case of SIO, the trajectories τ_1, \dots, τ_K and the energies e_1, \dots, e_K are fixed, so the state space is simply \mathcal{N}^K . In the case of BOO, the irradiation geometry is usually partitioned into subsets $\mathcal{T}_1, \dots, \mathcal{T}_L$ defining equispaced beam orientations and having the same size $J^{-1}K$, where J is the number of orientations imposed for the treatment; the state space is then

$$\Omega_{\text{BOO}} := C_J([1 \dots L]) \times \mathcal{N}^K, \quad (2.4)$$

where $C_J([1 \dots L])$ is the set of J -combinations of $[1 \dots L]$. Therefore,

$$\frac{|\Omega|}{|\Omega_{\text{BOO}}|} = \binom{L}{J}^{-1} \left(\sum_{\tau \in \mathcal{T}} |\mathcal{E}(\tau)| \right)^K. \quad (2.5)$$

In charged-particle therapy, J is typically 2–4, L does not exceed 72 (which corresponds to a step angle of 5° over a full circular orbit), and $|\mathcal{T}|$ is at least of the order of 10^4 . It follows that

$$\frac{|\Omega|}{|\Omega_{\text{BOO}}|} \gg J! |\mathcal{T}|^{K-J} \overline{|\mathcal{E}|}^K, \quad (2.6)$$

where $\overline{|\mathcal{E}|}$ is the average number of energies per trajectory:

$$\overline{|\mathcal{E}|} := |\mathcal{T}|^{-1} \sum_{\tau \in \mathcal{T}} |\mathcal{E}(\tau)|. \quad (2.7)$$

In other words, the size ratio of Ω to Ω_{BOO} grows exponentially with the number of beams in the treatment plan.

2.2 Objective function

We define an optimal treatment plan to be a global minimum of an objective function $U : \Omega \rightarrow \mathbb{R}$ measuring the conformity between the prescribed dose, or *dose plan*, and the dose

delivered by the input treatment plan. So the full ITP problem is $\min_{\omega} U(\omega)$ subject to $\omega := (\mathbf{b}_1, \dots, \mathbf{b}_K) \in \Omega$, or, equivalently,

$$\begin{aligned} & \min_{(\tau_k, e_k, n_k) : k \in [1..K]} U((\tau_1, e_1, n_1), \dots, (\tau_K, e_K, n_K)) \\ & \text{subject to } (\tau_k, e_k, n_k) \in \mathcal{T} \times \mathcal{E}(\tau_k) \times \mathcal{N} \text{ for all } k. \end{aligned} \tag{2.8}$$

We consider the standard objective defined by discretizing the PTV and the OAR region into a set of voxels \mathcal{V} and summing the squared distances between the prescribed and delivered doses over \mathcal{V} , that is,

$$U(\omega) := \sum_{v \in \mathcal{V}} \left(D_{\star}(v) - \sum_{k \in [1..K]} D(v, \mathbf{b}_k) \right)^2, \tag{2.9}$$

where D_{\star} is the dose plan and $D(v, \mathbf{b}_k)$ is the physical dose deposited in v by the k th beam of ω . The dose plan is, of course, zero in the OAR region, and the computation of the deposited doses is described in Sect. 5.1. Note that since the fluence is nonnegative, the state space implicitly imposes a positivity constraint on the delivered dose.

It is important to realize that the relation between the objective and the beam parameters cannot be expressed explicitly, because the deposited doses depend on the depth-dose profiles of the irradiating particles (these profiles are tabulated functions of the penetration depth that are parameterized by the energy and obtained experimentally or simulated numerically). This, along with the large size of the state space, makes full ITP particularly challenging.²

The BOO methods proposed in [31–33] search for a few optimal orientations in which to distribute the beam trajectories; the number of these active orientations is set manually [31,32] or controlled by a regularization parameter [33]. In contrast, full ITP searches for an optimal set of beams with no arrangement preference. In particular, if \mathcal{T} is a BOO geometry, full ITP will produce treatment plans whose beams are sparsely distributed across all orientations.³ In the case where \mathcal{T} is not partitioned into preferential orientations, full ITP should be viewed as a prospect for future charged-particle therapy systems.

3 Optimization by simulated annealing

This section provides a general description of SA under the usual assumption that the state space is finite, as is the case for full ITP. We refer to [58,59] for a comprehensive introduction to annealing-based optimization. Given an objective function $U : \Omega \rightarrow \mathbb{R}$ with $|\Omega| < \infty$, the goal is to find a state with the smallest possible value of U , ideally in the set of global minima

$$\Omega_{\min} := \{ \omega \in \Omega : U(\omega) = \min_{\Omega} U \}. \tag{3.1}$$

Here, Ω is *any* finite set (not necessarily a set of vectors as in Sect. 2), so we do not use the boldface notation for the states ω .

² In fact, the full ITP problem is an instance of the so-called minimum penalty treatment problem, which is APX-complete [60].

³ The number of active orientations could be controlled by adding a group-sparsity penalty to the objective, as in [33], but this topic is beyond the scope of this paper.

3.1 The simulated annealing algorithm

The main components of an SA algorithm are its *communication mechanism* and its *cooling schedule*. The communication mechanism defines the state exploration strategy and is represented by a Markov matrix on Ω , or, equivalently, a family $\Theta := (\Theta(\omega, \cdot))_{\omega \in \Omega}$ of probability distributions on Ω . The entries of Θ specify the possible moves in the state space, that is, $\Theta(\omega, \omega')$ is the probability of selecting ω' as the candidate solution given the current solution ω . Thinking of Ω as a set of points in \mathbb{R}^2 and of U as a measure of altitude, we call the triplet (Ω, Θ, U) a *landscape*. The cooling schedule is a decreasing sequence of temperatures $(T_m)_{m \in \mathbb{N}^*}$ that converges to zero and controls the acceptance of uphill moves. More precisely, the probability of accepting a move from ω to ω' at temperature T is given by

$$A_T(\omega, \omega') := \begin{cases} 1 & \text{if } U(\omega') \leq U(\omega), \\ \exp(-T^{-1}(U(\omega') - U(\omega))) & \text{if } U(\omega') > U(\omega). \end{cases} \quad (3.2)$$

Therefore, while downhill moves are unconditionally accepted, the probability of moving uphill decreases with decreasing temperature and vanishes as the temperature goes to zero.

A finite-time realization $(\omega_m)_{m \in [0..M]}$ of an SA algorithm with landscape (Ω, Θ, U) and cooling schedule $(T_m)_m$ is generated as follows:

```

choose an initial state  $\omega_0 \in \Omega$ 
for  $m = 1$  to  $M$  do
  draw a state  $\omega'$  from the probability distribution  $\Theta(\omega_{m-1}, \cdot)$ 
  set  $\Delta U \leftarrow U(\omega') - U(\omega_{m-1})$ 
  if  $\Delta U \leq 0$  then set  $\omega_m \leftarrow \omega'$ 
  else
    choose  $\kappa \in [0, 1]$  u.a.r.
    if  $\kappa \leq \exp(-T_m^{-1} \Delta U)$  then set  $\omega_m \leftarrow \omega'$ 
    else set  $\omega_m \leftarrow \omega_{m-1}$ 
  end(if)
end(if)
end(for)

```

3.2 The communication mechanism

The convergence of SA is based on the assumptions that the communication mechanism is *irreducible* and has *symmetric support*. More precisely:

- (i) there is a *communication path* from any state ω to any other state ω' , that is, a finite sequence of states $(\omega_j)_{j \in [0..p]}$ such that $\omega_0 = \omega$, $\omega_p = \omega'$, and $\Theta(\omega_{j-1}, \omega_j) > 0$ for all $j \in [1..p]$;
- (ii) $\Theta(\omega, \omega') > 0 \implies \Theta(\omega', \omega) > 0$ for all $(\omega, \omega') \in \Omega^2$.

Put simply, irreducibility means that the state space can be fully explored starting from any state, and symmetry guarantees that the SA algorithm can move backward. We also define the *communication distance* between two distinct states as the minimum length of the communication paths connecting them, and we say that two states *communicate at height h* if they are connected by a communication path along which the objective does not exceed h .

The set of allowed moves is often defined via a *neighborhood system* on Ω , that is, a collection $\{S(\omega)\}_{\omega \in \Omega}$ of subsets of Ω such that

- (i) $\omega \notin \mathcal{S}(\omega)$ for all $\omega \in \Omega$,
- (ii) $\omega' \in \mathcal{S}(\omega) \implies \omega \in \mathcal{S}(\omega')$ for all $(\omega, \omega') \in \Omega^2$.

Given such a system, a communication matrix Θ is constructed by assigning a weight $\theta(\omega, \omega') > 0$ to each ordered pair (ω, ω') of neighboring states and setting

$$\Theta(\omega, \omega') := \begin{cases} c \theta(\omega, \omega') & \text{if } \omega' \in \mathcal{S}(\omega), \\ 0 & \text{if } \omega' \notin \mathcal{S}(\omega) \cup \{\omega\}, \\ 1 - c \sum_{\tilde{\omega} \in \mathcal{S}(\omega)} \theta(\omega, \tilde{\omega}) & \text{if } \omega' = \omega, \end{cases} \tag{3.3}$$

where the constant c is chosen small enough so that Θ is a Markov matrix:

$$\max_{\omega \in \Omega} \sum_{\tilde{\omega} \in \mathcal{S}(\omega)} \theta(\omega, \tilde{\omega}) \leq \frac{1}{c}. \tag{3.4}$$

The irreducibility is then equivalent to having a neighbor-to-neighbor path between any two distinct states, and the support is symmetric by construction.

The communication mechanism also defines the set of local minima:

$$\Omega_{\text{loc}}(\Theta) := \{ \omega \in \Omega \mid \forall \omega' \in \Omega, \Theta(\omega, \omega') > 0 \implies U(\omega) \leq U(\omega') \} \tag{3.5}$$

[in other words, $\omega \in \Omega_{\text{loc}}(\Theta)$ if and only if the objective is lower bounded by $U(\omega)$ in the neighborhood of ω]. A state ω in $\Omega_{\text{loc}}(\Theta) \setminus \Omega_{\text{min}}$ is called a *nonglobal minimum* and is often characterized by its *depth* and its *basin*: the depth of ω is the smallest number $d \geq 0$ such that ω communicates at height $U(\omega) + d$ with a state ω' satisfying $U(\omega') < U(\omega)$; the basin of ω is the set of states communicating with ω at a height less than $U(\omega) + d$.

3.3 Convergence properties

Formally, an SA algorithm with landscape (Ω, Θ, U) and cooling schedule $(T_m)_{m \in \mathbb{N}^*}$ is a discrete-time nonhomogeneous Markov chain $(X_m)_{m \in \mathbb{N}}$ with transitions

$$P(X_m = \omega' \mid X_{m-1} = \omega) = Q_{T_m}(\omega, \omega'), \tag{3.6}$$

where the family of Markov matrices $(Q_T)_{T > 0}$ is defined by

$$Q_T(\omega, \omega') := \begin{cases} \Theta(\omega, \omega') A_T(\omega, \omega') & \text{if } \omega' \neq \omega, \\ 1 - \sum_{\tilde{\omega} \in \Omega \setminus \{\omega\}} Q_T(\omega, \tilde{\omega}) & \text{if } \omega' = \omega. \end{cases} \tag{3.7}$$

The theoretical intuition behind SA is best understood when, in addition to being irreducible, the communication matrix Θ is symmetric. In this case, the stationary distribution of Q_T is the Gibbs distribution

$$\pi_T(\omega) := (Z_T)^{-1} \exp(-T^{-1}U(\omega)), \quad Z_T := \sum_{\omega \in \Omega} \exp(-T^{-1}U(\omega)). \tag{3.8}$$

On the one hand, π_T tends to the uniform distribution on Ω as $T \rightarrow +\infty$, and on the other hand, π_T tends to the uniform distribution on Ω_{min} as $T \rightarrow 0$. Therefore the state space is more freely explored at the beginning of the annealing process (i.e., when the temperature is high), and we can expect to get arbitrarily close to a global minimum in the long run (i.e., when the temperature approaches zero).

It is well-known that if $(T_m)_m$ is a logarithmic cooling schedule of the form

$$T_m = \frac{T_{\max} \log 2}{\log(m + 1)} \tag{3.9}$$

with T_{\max} greater than a critical value depending on the landscape, then SA converges to the set of global minima in the sense that

$$\lim_{m \rightarrow \infty} P(X_m \in \Omega_{\min}) = 1 \tag{3.10}$$

(see, e.g., [61,62]). However, this result holds for values of T_{\max} that are usually too large to reach the low-temperature regime in a reasonable amount of computation time. In fact, successful applications of SA use exponential cooling, the justification of which dates back to the early 1990's [63]. Two results are particularly important. First, the convergence rate of SA cannot be faster than some optimal power μ^* of m^{-1} . Second, for every $\mu \in (0, \mu^*)$, there is a family $\{(T_{m,M})_{m \in [1..M]}\}_{M \in \mathbb{N}^*}$ of finite exponential cooling schedules such that the final state X_M of the SA chain controlled by $(T_{m,M})_m$ satisfies

$$P(X_M \in \Omega_{\min}) \geq 1 - M^{-\mu} \text{ for sufficiently large } M. \tag{3.11}$$

The optimal convergence-speed exponent μ^* is the inverse of the so-called *difficulty* of the landscape, defined as the maximum ratio of the depth of nonglobal minima to their level above $\min_{\Omega} U$. The families of cooling schedules that yield (3.11) have the form

$$T_{m,M} := T_{\max} \left(\frac{T_{\min}}{T_{\max}} \right)^{(m-1)/(M-1)}, \tag{3.12}$$

where T_{\max} is independent of the horizon M and $\log T_{\min} \sim -\log M$.

3.4 Simulated annealing with temperature-dependent communication

Ideally, the communication mechanism should allow rapid exploration of the state space at high temperatures and maintain a reasonable acceptance rate at low temperatures. However, these two demands respectively require large and small neighborhoods $\mathcal{S}(\omega)$ in (3.3) and are thus in conflict. Consequently, practical SA algorithms are usually slowly mixing over wide temperature ranges. A simple way to overcome this limitation is to replace Θ in (3.7) by a family of Markov matrices $(\Theta_T)_{T>0}$ converging to a communication mechanism adapted to the low-temperature regime. This variant of SA belongs to the class of stochastic continuation algorithms [59,64] and shares the convergence properties of standard SA if the following weak conditions are satisfied:

- (i) $\lim_{T \rightarrow 0} \Theta_T(\omega, \omega')$ exists for all $(\omega, \omega') \in \Omega^2$;
- (ii) the matrix with entries $\Theta_0(\omega, \omega') := \lim_{T \rightarrow 0} \Theta_T(\omega, \omega')$ is an irreducible Markov matrix with symmetric support;
- (iii) the set of possible moves $\{(\omega, \omega') \in \Omega^2 : \Theta_T(\omega, \omega') > 0\}$ freezes to the support of Θ_0 at low temperatures.

In this case, the convergence-speed exponent of SA with temperature-dependent communication can be made arbitrarily close to the optimal exponent μ^* of standard SA [59]. More precisely, for every $\mu \in (0, \mu^*)$, (3.11) holds for piecewise-constant exponential cooling schedules of the form

$$T_{m,M} := T_{\max} \left(\frac{T_{\min}}{T_{\max}} \right)^{\alpha_{m,M}}, \quad \alpha_{m,M} := (\lceil M^{-1} m \xi \rceil - 1) / (\xi - 1), \tag{3.13}$$

where ξ is the number of constant-temperature stages.

4 Design of the communication mechanism

We present here communication mechanisms for full ITP in increasing order of complexity and efficiency. This incremental design approach facilitates the presentation and justification of our final strategy for exploring the state space. Each proposed mechanism generates a candidate move from the current solution by changing the parameters of a randomly selected beam b_ℓ ; we call the resulting new beam b'_ℓ a *candidate beam*.

We begin with *type-1 moves* which generate candidate beams by choosing their trajectories, energies and fluences u.a.r. Although SA with type-1 moves allows to demonstrate the feasibility of full ITP, it is slowly mixing in the mid- and low-temperature regimes and thus inefficient. To avoid this problem, we introduce *type-2 moves* which select the trajectories and energies u.a.r. and set the fluences to optimal values. We continue further with *type-3 moves* which select the trajectories u.a.r. and find optimal energy-fluence pairs in randomly chosen partition blocks of the sets $\mathcal{E}(\tau) \times \mathcal{N}$, $\tau \in \mathcal{T}$. Our final exploration strategy consists of temperature-dependent type-3 moves along with occasional type-1 moves to ensure irreducibility and symmetric support.

The following sections provide high-level descriptions of these increasingly sophisticated communication strategies; implementation details and computational cost analyses are given in Sect. 5.

4.1 Type-1 moves (uniform candidate beams)

A basic mechanism to generate a candidate solution $\omega' := (b'_1, \dots, b'_K)$ from $\omega := (b_1, \dots, b_K)$ is to pick a beam index ℓ u.a.r. and then change the trajectory, energy and fluence of b_ℓ u.a.r. This is implemented as follows:

- choose $\ell \in [1 \dots K]$ u.a.r.
- set $b'_k \leftarrow b_k$ for all $k \neq \ell$
- choose $\tau \in \mathcal{T}$ u.a.r.
- choose $(e, n) \in \mathcal{E}(\tau) \times \mathcal{N}$ u.a.r.
- set $b'_\ell \leftarrow (\tau, e, n)$

The corresponding communication matrix is denoted by Θ_1 and defined by

$$\Theta_1(\omega, \omega') := \begin{cases} \frac{K^{-1}}{|\mathcal{T}| \cdot |\mathcal{E}(\tau)| \cdot |\mathcal{N}|} & \text{if } \omega' \in S_1(\omega), \\ 0 & \text{if } \omega' \notin S_1(\omega) \cup \{\omega\}, \end{cases} \tag{4.1}$$

where $S_1(\omega)$ is the set of feasible treatment plans that differ from ω by a single beam b'_ℓ , that is,

$$\omega' \in S_1(\omega) \iff \omega' \in \Omega \text{ and } \exists! \ell \in [1 \dots K], b'_\ell \neq b_\ell, \tag{4.2}$$

and where τ is the trajectory of b'_ℓ .⁴ We call a move from ω to $\omega' \in S_1(\omega)$ a *type-1 move*. Clearly, $\{S_1(\omega)\}_{\omega \in \Omega}$ is a neighborhood system and any two states in Ω are connected by a sequence of at most K type-1 moves. So Θ_1 is irreducible and has symmetric support.

The problem with type-1 moves is that their rejection rate is too high in the low-temperature regime, causing slow convergence of the SA chain. Indeed, Θ_1 generates candidate beams by sampling from a large set of three-component vectors [namely, $(\mathcal{T} \otimes \mathcal{E}) \times \mathcal{N}$], which often

⁴ Note that Θ_1 has the form (3.3) with weights $\theta(\omega, \omega') = |\mathcal{E}(\tau)|^{-1}$.

induces significant changes in the objective and hence a very low probability to move uphill at low temperatures.

4.2 Type-2 moves (optimal fluence moves)

To improve the mixing rate compared with type-1 moves, a natural idea is to reduce the beam sample-space while encouraging moves toward minima. A first step in this direction is to set the fluence of the candidate beam $\mathbf{b}'_\ell := (\tau, e, n)$ so as to obtain the smallest possible objective value when the trajectory τ and the energy e are fixed. That is, we set n equal to

$$n_{\omega \setminus \ell}(\tau, e) := \arg \min_{\nu \in \mathcal{N}} U(\omega \setminus \ell(\tau, e, \nu)), \tag{4.3}$$

where $\omega \setminus \ell(\tau, e, \nu)$ denotes the treatment plan obtained by replacing the beam \mathbf{b}_ℓ of the current state ω by (τ, e, ν) .⁵ This choice is interesting because the function

$$x \in \mathbb{R} \mapsto U(\omega \setminus \ell(\tau, e, x)) \tag{4.4}$$

is a polynomial of degree 2 whose coefficients can be computed with negligible overhead (details for computing the optimal fluence $n_{\omega \setminus \ell}(\tau, e)$ are given in Sect. 5.3). The corresponding communication matrix, denoted by Θ_2 , is defined by

$$\Theta_2(\omega, \omega') := \begin{cases} K^{-1} & \text{if } \omega' \in \mathcal{S}_2(\omega), \\ \frac{|T| \cdot |\mathcal{E}(\tau)|}{|T| \cdot |\mathcal{E}(\tau)|} & \text{if } \omega' \notin \mathcal{S}_2(\omega) \cup \{\omega\}, \\ 0 & \end{cases} \tag{4.5}$$

where $\mathcal{S}_2(\omega)$ is the set of feasible treatment plans that differ from ω by a single beam with optimal fluence:

$$\omega' \in \mathcal{S}_2(\omega) \iff \begin{cases} \omega' \in \mathcal{S}_1(\omega), \\ (\forall k)[\mathbf{b}'_k = \mathbf{b}_k \text{ or } n'_k = n_{\omega \setminus k}(\tau'_k, e'_k)] \end{cases} \tag{4.6}$$

(τ'_k, e'_k , and n'_k are the trajectory, energy and fluence of \mathbf{b}'_k , the k th beam of ω'). We call a move from ω to $\omega' \in \mathcal{S}_2(\omega)$ a *type-2 move*.

Despite their appeal, type-2 moves cannot be used alone because Θ_2 is not irreducible and does not have symmetric support; indeed:

- (i) there exist inaccessible states, simple examples of which are the treatment plans $\check{\omega}(n) := ((\tau, e, 0), \dots, (\tau, e, 0), (\tau, e, n))$ whose beams have the same trajectory and energy and are all inactive except for one with a fluence n smaller than $\arg \min_{\nu \in \mathcal{N}} U(\check{\omega}(\nu))$;
- (ii) the implication $\omega' \in \mathcal{S}_2(\omega) \implies \omega \in \mathcal{S}_2(\omega')$ holds for all ω' only if every beam \mathbf{b}_k of ω is such that $n_k = n_{\omega \setminus k}(\tau_k, e_k)$.⁶

A way around this problem is to allow occasional type-1 moves using the mixed communication mechanism

$$\Theta_{2,\eta} := (1 - \eta)\Theta_2 + \eta\Theta_1, \tag{4.7}$$

⁵ Although the possibility cannot be excluded, the argument of the minimum in (4.3) is unlikely to contain more than one element; so we assume it is a singleton for simplicity.

⁶ Therefore $\{\mathcal{S}_2(\omega)\}_{\omega \in \Omega}$ is not a neighborhood system.

where $\eta \in (0, 1)$ is the probability of choosing a type-1 rather than a type-2 move. Since $S_2(\omega) \subset S_1(\omega)$ for all ω , we have $\Theta_{2,\eta}(\omega, \omega') > 0 \iff \omega' \in S_1(\omega)$ whenever $\omega \neq \omega'$, so $\Theta_{2,\eta}$ inherits the irreducibility and symmetry of Θ_1 . In practice, a state ω' is drawn from the probability distribution $\Theta_{2,\eta}(\omega, \cdot)$ as follows:

```

choose  $\ell \in [1..K]$  u.a.r.
set  $\mathbf{b}'_k \leftarrow \mathbf{b}_k$  for all  $k \neq \ell$ 
choose  $\tau \in \mathcal{T}$ ,  $e \in \mathcal{E}(\tau)$  and  $\kappa \in [0, 1]$  u.a.r.
if  $\kappa < \eta$  then choose  $n \in \mathcal{N}$  u.a.r.
else set  $n \leftarrow n_{\omega_\ell}(\tau, e)$ 
end(if)
set  $\mathbf{b}'_\ell \leftarrow (\tau, e, n)$ 
    
```

Obviously, η should be chosen small enough to fully benefit from type-2 moves (e.g., $\eta = 0.01$, as in our experiments).

4.3 Type-3 moves (semi-optimal energy-fluence moves)

As an extension of type-2 moves, we can think of further reducing the size of the beam sample-space by simultaneously optimizing the energy and fluence. The problem to be solved at each iteration is then the following: given an index ℓ and a trajectory τ , find an energy-fluence pair (e, n) such that

$$\begin{aligned}
 U(\omega_{\ell}(\tau, e, n)) &= \min_{(\epsilon, v) \in \mathcal{E}(\tau) \times \mathcal{N}} U(\omega_{\ell}(\tau, \epsilon, v)) \\
 &= \min_{\epsilon \in \mathcal{E}(\tau)} U(\omega_{\ell}(\tau, \epsilon, n_{\omega_\ell}(\tau, \epsilon))).
 \end{aligned}
 \tag{4.8}$$

Because of the intricate relationship between deposited dose and beam energy, this task requires computing the whole set of optimal fluences $\{n_{\omega_\ell}(\tau, \epsilon)\}_{\epsilon \in \mathcal{E}(\tau)}$. Therefore, even though $n_{\omega_\ell}(\tau, \epsilon)$ is fast to obtain, the cost of computing the optimal energy-fluence pair can become prohibitive as the energy increment δ_e [see (2.1)] decreases. Setting a lower bound on δ_e to limit the computational burden is not a viable option, since it affects the precision of Bragg peak positioning and hence the conformity to the dose plan. We propose instead to restrict the set of possible energies $\mathcal{E}(\tau)$ in (4.8) to subsets of the form

$$\mathcal{E}_{p,j}(\tau) := \{p(l + j)\delta_e\}_{l \in \mathbb{N}} \cap [e_{\min}(\tau) - \delta_e, e_{\max}(\tau) + \delta_e],
 \tag{4.9}$$

where $p \in [1..|\mathcal{E}(\tau)|]$ and j is the realization of a uniform discrete random variable on $[0..p-1]$. The collection $\{\mathcal{E}_{p,j}(\tau)\}_{j \in [0..p-1]}$ is a partition of $\mathcal{E}(\tau)$ into p interlaced subsets of equispaced energies, each of size $\lfloor p^{-1}|\mathcal{E}(\tau)| \rfloor$ or $\lceil p^{-1}|\mathcal{E}(\tau)| \rceil$. So on the one hand, the computational cost is controlled by the parameter p , and on the other hand, the precision of Bragg peak positioning is not affected because no energy value is discarded.

Let $p \in \mathbb{N}^*$. Given the current state ω , we say that $\mathbf{b}'_\ell := (\tau, e, n)$ is a *type-3 candidate beam* if (e, n) minimizes the partial objective $(\epsilon, v) \mapsto U(\omega_{\ell}(\tau, \epsilon, v))$ over $\mathcal{E}_{p,j}(\tau) \times \mathcal{N}$ for some $j \in [0..p-1]$. Let $S_{3,p}(\omega)$ be the set of treatment plans that differ from ω by a single type-3 candidate beam. Using the definition of S_2 in (4.6), we have

$$\omega' \in S_{3,p}(\omega) \iff \begin{cases} \omega' \in S_2(\omega), \\ (\forall k)[\mathbf{b}'_k = \mathbf{b}_k \text{ or } e'_k \in \{e_{\omega_{k,p,j}}(\tau'_k)\}_{j \in [0..p-1]}], \end{cases}
 \tag{4.10}$$

where $e_{\omega_{\setminus k}, p, j}(\tau)$ denotes the energy of the best candidate beam in the sample space $\{\tau\} \times \mathcal{E}_{p, j}(\tau) \times \mathcal{N}$:⁷

$$e_{\omega_{\setminus k}, p, j}(\tau) := \arg \min_{\epsilon \in \mathcal{E}_{p, j}(\tau)} U(\omega_{\setminus k}(\tau, \epsilon, n_{\omega_{\setminus k}}(\tau, \epsilon))). \tag{4.11}$$

We call a move from ω to $\omega' \in \mathcal{S}_{3, p}(\omega)$ a *type-3 move*; the associated communication matrix, say $\Theta_{3, p}$, has the same form as Θ_2 , but with smaller support:

$$\Theta_{3, p}(\omega, \omega') := \begin{cases} \frac{K^{-1}}{|T|p} & \text{if } \omega' \in \mathcal{S}_{3, p}(\omega), \\ 0 & \text{if } \omega' \notin \mathcal{S}_{3, p}(\omega) \cup \{\omega\}. \end{cases} \tag{4.12}$$

As for type-2 moves, $\Theta_{3, p}$ is neither irreducible nor has symmetric support, a problem overcome by allowing occasional type-1 moves. This leads to the mixed communication mechanism

$$\Theta_{3, p, \eta} := (1 - \eta)\Theta_{3, p} + \eta\Theta_1, \quad \eta \in (0, 1), \tag{4.13}$$

which is implemented as follows:

```

choose  $\ell \in [1 \dots K]$  u.a.r.
set  $b'_k \leftarrow b_k$  for all  $k \neq \ell$ 
choose  $\tau \in T$  and  $\kappa \in [0, 1]$  u.a.r.
if  $\kappa < \eta$  then
    choose  $e \in \mathcal{E}(\tau)$  and  $n \in \mathcal{N}$  u.a.r.
else
    choose  $j \in [0 \dots p - 1]$  u.a.r.
    compute  $n_{\omega_{\setminus \ell}}(\tau, \epsilon)$  for every  $\epsilon \in \mathcal{E}_{p, j}(\tau)$ 
    set  $e \leftarrow e_{\omega_{\setminus \ell}, p, j}(\tau)$ 
    set  $n \leftarrow n_{\omega_{\setminus \ell}}(\tau, e)$ 
end(if)
set  $b'_\ell \leftarrow (\tau, e, n)$ 
    
```

The use of type-3 moves should be viewed as a speed-up technique rather than a heuristic. Indeed, the mixed communication mechanism $\Theta_{3, p, \eta}$ preserves the global convergence properties of SA and does not discard any energy value; so neither optimality nor accuracy is traded for speed. We could speak of a heuristic if we restricted $\mathcal{E}(\tau)$ to a subset $\mathcal{E}_{p, j}(\tau)$ with a fixed j , because it would decrease the cost of computing a type-3 candidate beam by a factor of p but would also degrade the depth resolution by the same factor. Instead, we pick j u.a.r. so as to consider all the energy partition blocks $\mathcal{E}_{p, j}(\tau)$. In this way, type-3 candidate beams have optimal energy and fluence with probability p^{-1} , and so depth resolution is not degraded.

The smaller p , the greater the tendency to move downhill. However, it is important to emphasize that type-3 moves are not necessarily downhill, even when $p = 1$. The fact that a type-3 move from ω to $\omega' := \omega_{\setminus \ell}(\tau, e, n)$ is uphill can be written as

$$U(\omega_{\setminus \ell}) - U(\omega) > U(\omega_{\setminus \ell}) - U(\omega'), \tag{4.14}$$

where $\omega_{\setminus \ell}$ denotes the treatment plan obtained by deactivating the ℓ th beam (that is, by setting its fluence to zero); in other words, the variation in the objective resulting from deactivating the ℓ th beam is larger for ω than for ω' . This occurs, for example, when the following conditions are met:

⁷ Again, for simplicity, we assume that the argument of the minimum is a singleton.

- (i) the beam b_ℓ in ω does not intersect any OAR, and the dose delivered by ω is below the prescribed dose in all the PTV voxels intersected by b_ℓ [so $U(\omega_{\setminus \ell}) - U(\omega) > 0$];
- (ii) the type-3 candidate beam b'_ℓ enters through an OAR, and the dose delivered by $\omega_{\setminus \ell}$ is above the prescribed dose in all the PTV voxels intersected by b'_ℓ [in which case the optimal fluence is zero, so $U(\omega_{\setminus \ell}) - U(\omega') = 0$].

4.4 Dynamic exploration strategy

The number p of energy partition blocks in type-3 moves balances the mixing rate and the computational cost: the smaller p , the higher the acceptance rate, but the more expensive the iterations. The optimal value of p , however, depends on the irradiation geometry and can only be found empirically. To avoid having to choose p , we propose using a dynamic mechanism in which p decreases with decreasing temperature—the rationale is that making efforts to find optimal candidate beams in the high-temperature regime is a waste of time, since most moves are accepted due to the high acceptance rate. Conversely, encouraging moves toward local minima in the low-temperature regime maintains a reasonable acceptance rate and hence increases the performance of SA.

Our final strategy for exploring the state space consists of the mechanism (4.13) with p decreasing at each temperature stage of a piecewise-constant cooling schedule of the form (3.13). More specifically, we use a sequence of communication matrices

$$(\Theta_{3,p_m,\eta})_{m \in [1..M]}, \quad p_m := p_{\max} - \lceil (\alpha_{m,M})^\beta (p_{\max} - 1) \rceil, \tag{4.15}$$

where p_{\max} is the initial number of energy partition blocks, $\lceil \cdot \rceil$ is the nearest integer function, and the exponent $\beta > 0$ controls the decreasing rate of p_m .⁸ An important point is that the performance of SA with the communication sequence $(\Theta_{3,p_m,\eta})_m$ is not sensitive to the choice of p_{\max} and β . Indeed, we can fix the values of these parameters regardless of the irradiation geometry and achieve performance similar to that obtained using the static mechanism $\Theta_{3,p,\eta}$ with an optimal (geometry-dependent) value of p .

At the beginning of the annealing process, type-3 moves should involve only a few energies so that the iterations are cheap. So we set p_{\max} to half the average number of energies per trajectory, that is,

$$p_{\max} := \left\lceil \frac{1}{2} |\mathcal{E}| \right\rceil. \tag{4.16}$$

On the other extreme, $p_m = 1$ in the last temperature stage (i.e., when $1 + \xi^{-1} < M^{-1}m \leq 1$), so the annealing process ends with optimal candidate beams at the price of expensive iterations. Therefore the exponent β balances speed and accuracy: the smaller β , the faster p_m decreases, and thus the higher the computational cost and the expected quality of the solution. According to our experience, $\beta = 0.2$ is a consistently good compromise.

5 Implementation details and computational cost

The computational complexity of SA is dominated by the evaluation of the variations in the objective caused by the moves in the search space. Here we focus on this critical task for the objective defined in Sect. 2 and the communication mechanisms described in Sect. 4. For clarity and completeness, we begin with a brief description of the physical dose model,

⁸ The sequence $(p_m)_{m \in [1..M]}$ is also piecewise constant, with a number of plateaus less than or equal to the number of temperature stages.

which is similar to the one proposed in [12]. We then follow the same steps as in Sect. 4: each analysis is built upon the previous one, starting with type-1 moves and concluding with the proposed dynamic exploration strategy.

5.1 Simulation of the physical dose

To compute the deposited doses $D(v, \mathbf{b}_k)$, we consider particle beams with Gaussian fluence profile and negligible lateral scattering.⁹ A beam \mathbf{b} is modeled by a finite set $\{\mathbf{b}(i)\}_{i \in \mathcal{I}}$ of parallel rays whose positions are obtained by sampling the transverse plane of \mathbf{b} using concentric mapping. Let σ be the standard deviation of the fluence profile, usually specified via the full width at half maximum $\text{FWHM} = 2\sigma(2 \log 2)^{1/2}$. The fluence contributions $\gamma(\mathbf{b}(i))$ of the rays are defined by

$$\begin{cases} \gamma(\mathbf{b}(i)) \propto \exp\left(-\frac{1}{2}(\rho(\mathbf{b}(i))/\sigma)^2\right), \\ \sum_{i \in \mathcal{I}} \gamma(\mathbf{b}(i)) = 1, \end{cases} \quad (5.1)$$

where $\rho(\mathbf{b}(i))$ is the radial position of $\mathbf{b}(i)$ with respect to the trajectory of \mathbf{b} . For every voxel v and every ray $\mathbf{b}(i)$ intersecting v , we denote by $d_{\text{in}}[v, \mathbf{b}(i)]$ and $d_{\text{out}}[v, \mathbf{b}(i)]$ the water-equivalent penetration depths corresponding to the entry and exit points of $\mathbf{b}(i)$ into and out of v , respectively. The dose deposited by $\mathbf{b} := (\tau, e, n)$ in v depends linearly on the fluence and nonlinearly on the energy:

$$D(v, \mathbf{b}) \propto \sum_{i \in \mathcal{I} : \mathbf{b}(i) \cap v \neq \emptyset} n \gamma(\mathbf{b}(i)) (E(e, d_{\text{out}}[v, \mathbf{b}(i)]) - E(e, d_{\text{in}}[v, \mathbf{b}(i)])), \quad (5.2)$$

where $E(e, \cdot)$ is the cumulative depth-dose profile for the energy e [that is, $E(e, d)$ is the energy loss at depth d for particles with energy e] and the proportionality factor depends on the mass of tissue in v .

At energies below or about 150 MeV u^{-1} , the depth-dose profiles $\partial E(e, \cdot)/\partial d$ have very sharp Bragg peaks, making it more difficult to deliver a smooth dose distribution at small depths. A simple and effective way to avoid this problem is to smooth the longitudinal dose distribution using a passive, static beam-shaping element called a *ripple filter* [65]. The use of a ripple filter is easily accounted for by replacing the cumulative depth-dose profile $E(e, \cdot)$ in (5.2) by the smoothed version

$$\tilde{E}(e, d) := a^{-1} \int_0^a E(e, d + f(x)) dx, \quad (5.3)$$

where the function $f : [0, a] \rightarrow \mathbb{R}_+$ represents the periodic groove profile of the filter (a is the half-period of this profile). Figure 1 illustrates the effect of a ripple filter on the dose deposited in water by carbon-ion beams.

The cumulative depth-dose profiles are precomputed by Geant4 Monte Carlo simulations [66] and stored in memory after being transformed as in (5.3). Therefore, the major part of the computation time for estimating the deposited doses is consumed in calculating ray-voxel intersections. This can be avoided by storing the set of fluence-normalized doses deposited in all voxels for all trajectories and energies, that is, the set

$$\mathcal{D} := \{D(v, (\tau, e, 1)) : v \in \mathcal{V}, (\tau, e) \in \mathcal{T} \otimes \mathcal{E}\}. \quad (5.4)$$

⁹ This approximation is justified for charged particles, with a greater accuracy for heavy ions such as carbon than for light ions such as protons.

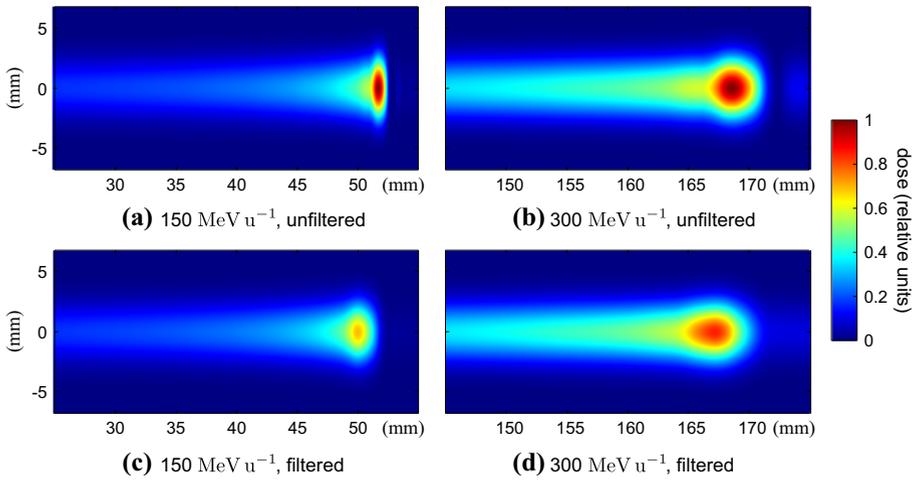


Fig. 1 Effect of a 2 mm ripple filter on the dose deposited in water by ^{12}C -ion beams with a 4 mm FWHM: **a**, **b** 150 and 300 MeV u^{-1} unfiltered beams; **c**, **d** 150 and 300 MeV u^{-1} beams modulated by the ripple filter

In the computational cost analyses given in the following sections, we distinguish between the case where \mathcal{D} is stored in memory and the case where it is not because of insufficient capacity.

5.2 Type-1 moves

Let $\mathcal{V}(\mathbf{b})$ denote the set of voxels in \mathcal{V} intersected by a beam \mathbf{b} (this set is defined by the trajectory and the radius of \mathbf{b}), and let $\mathbf{b}'_\ell := (\tau, e, n)$ be a candidate beam defining a move from ω to $\omega' \in \mathcal{S}_1(\omega)$. Then

$$U(\omega') - U(\omega) = \sum_{v \in \mathcal{V}(\mathbf{b}_\ell) \cup \mathcal{V}(\mathbf{b}'_\ell)} \delta_D(v, \ell) (\delta_D(v, \ell) - \Delta_D(v)), \tag{5.5}$$

where

$$\delta_D(v, \ell) := D(v, \mathbf{b}'_\ell) - D(v, \mathbf{b}_\ell) \tag{5.6}$$

and

$$\Delta_D(v) := 2 \left(D_\star(v) - \sum_{k \in [1..K]} D(v, \mathbf{b}_k) \right). \tag{5.7}$$

This suggests to store the sets

$$\mathcal{D}(\omega) := \bigcup_{k \in [1..K]} \{D(v, \mathbf{b}_k)\}_{v \in \mathcal{V}(\mathbf{b}_k)} \tag{5.8}$$

and

$$\Delta \mathcal{D}(\omega) := \{\Delta_D(v)\}_{v \in \mathcal{V}} \tag{5.9}$$

and update them as follows: if the move from ω to ω' is accepted, then set

$$\Delta_D(v) \leftarrow \Delta_D(v) - 2\delta_D(v, \ell) \quad \text{for all } v \in \mathcal{V}(\mathbf{b}_\ell) \cup \mathcal{V}(\mathbf{b}'_\ell) \tag{5.10}$$

to first obtain $\Delta\mathcal{D}(\omega')$, and set

$$D(v, \mathbf{b}_\ell) \leftarrow D(v, \mathbf{b}'_\ell) \quad \text{for all } v \in \mathcal{V}(\mathbf{b}'_\ell) \tag{5.11}$$

to obtain $\mathcal{D}(\omega')$. In this way, given $\mathcal{D}(\omega)$, $\Delta\mathcal{D}(\omega)$, and the set of deposited doses

$$\mathcal{D}(\mathbf{b}'_\ell) := \{D(v, \mathbf{b}'_\ell)\}_{v \in \mathcal{V}(\mathbf{b}'_\ell)}, \tag{5.12}$$

the number of flops¹⁰ for computing the variation (5.5) is

$$N_{\Delta U}(\omega, \omega') := 2|\mathcal{V}(\mathbf{b}_\ell) \ominus \mathcal{V}(\mathbf{b}'_\ell)| + 3|\mathcal{V}(\mathbf{b}_\ell) \cap \mathcal{V}(\mathbf{b}'_\ell)|, \tag{5.13}$$

where \ominus is the symmetric difference operator [i.e., $\mathcal{A} \ominus \mathcal{B} := (\mathcal{A} \cup \mathcal{B}) \setminus (\mathcal{A} \cap \mathcal{B})$]. Furthermore, if the dose differences (5.6) are stored during the computation of (5.5), the number of flops for updating $\Delta\mathcal{D}(\omega)$ to $\Delta\mathcal{D}(\omega')$ is

$$N_{\Delta\mathcal{D}}(\omega, \omega') := 2|\mathcal{V}(\mathbf{b}_\ell) \cup \mathcal{V}(\mathbf{b}'_\ell)|. \tag{5.14}$$

The cost of computing the set of deposited doses $\mathcal{D}(\mathbf{b}'_\ell)$ depends on whether the set \mathcal{S} of fluence-normalized doses [see (5.4)] is stored in memory. If so, computing $\mathcal{D}(\mathbf{b}'_\ell)$ requires only $|\mathcal{V}(\mathbf{b}'_\ell)|$ multiplications,¹¹ and therefore the number of flops for implementing a type-1 move is

$$\begin{aligned} N_1(\omega, \omega') &:= N_{\Delta U}(\omega, \omega') + N_{\Delta\mathcal{D}}(\omega, \omega') + |\mathcal{V}(\mathbf{b}'_\ell)| \\ &= 4|\mathcal{V}(\mathbf{b}_\ell)| + 5|\mathcal{V}(\mathbf{b}'_\ell)| - 3|\mathcal{V}(\mathbf{b}_\ell) \cap \mathcal{V}(\mathbf{b}'_\ell)|. \end{aligned} \tag{5.15}$$

In contrast, if \mathcal{S} cannot be stored in memory, computing $\mathcal{D}(\mathbf{b}'_\ell)$ makes up the bulk of the load: the number of flops required in addition to $N_1(\omega, \omega')$ is

$$\begin{aligned} N_{\mathcal{S},1}(\omega, \omega') &\approx (3 + N_d) \sum_{v \in \mathcal{V}(\mathbf{b}'_\ell)} |\{i \in \mathcal{I} : \mathbf{b}'_\ell(i) \cap v \neq \emptyset\}| \\ &= (3 + N_d) \sum_{i \in \mathcal{I}} |\mathcal{V}(\mathbf{b}'_\ell(i))|, \end{aligned} \tag{5.16}$$

where N_d is the average number of flops for computing the entry and exit depths of a ray into a voxel, and where $\mathcal{V}(\mathbf{b}'_\ell(i))$ denotes the set of voxels intersected by the i th ray of \mathbf{b}'_ℓ .

5.3 Type-2 moves

Let τ and e be the trajectory and energy of a candidate beam \mathbf{b}'_ℓ defining a move from ω to $\omega' \in \mathcal{S}_2(\omega)$. The optimal fluence $n_{\omega \setminus \ell}(\tau, e)$ is obtained by finding the global minimum v^* of the function of a real variable (4.4) and then setting

$$n_{\omega \setminus \ell}(\tau, e) := \begin{cases} \lceil \delta_n^{-1} v^* \rceil \delta_n & \text{if } v^* \in [0, n_{\max}], \\ 0 & \text{if } v^* < 0, \\ n_{\max} & \text{if } v^* > n_{\max}. \end{cases} \tag{5.17}$$

Therefore the computational cost of implementing a type-2 move is essentially that of a type-1 move plus that of finding v^* . Let $\underline{\mathbf{b}} := (\tau, e, 1)$, so that $D(v, \underline{\mathbf{b}})$ is the fluence-normalized

¹⁰ Throughout the paper, *flops* is the plural of flop (an elementary floating point operation) and is not to be confused with “flops per second”.

¹¹ We recall that $D(v, \mathbf{b}'_\ell)$ is the product of the fluence of \mathbf{b}'_ℓ and the fluence-normalized dose $D(v, (\tau, e, 1))$.

dose deposited in v by \mathbf{b}'_ℓ . It is not difficult to see that v^* is the minimum of the second-order polynomial

$$x \in \mathbb{R} \mapsto \sum_{v \in \mathcal{V}(\mathbf{b}'_\ell)} \left(x^2 D^2(v, \mathbf{b}) - x D(v, \mathbf{b}) (2D(v, \mathbf{b}_\ell) + \Delta_D(v)) \right), \tag{5.18}$$

that is,

$$v^* = \left(\sum_{v \in \mathcal{V}(\mathbf{b}'_\ell)} D(v, \mathbf{b}) (2D(v, \mathbf{b}_\ell) + \Delta_D(v)) \right) \left(2 \sum_{v \in \mathcal{V}(\mathbf{b}'_\ell)} D^2(v, \mathbf{b}) \right)^{-1}. \tag{5.19}$$

Suppose the values of the denominator in this expression are precomputed and stored for all trajectory-energy pairs $(\tau, e) \in \mathcal{T} \otimes \mathcal{E}$. If \mathcal{D} is in memory, the computation of v^* requires $4|\mathcal{V}(\mathbf{b}'_\ell)|$ flops, which is roughly half the cost of a type-1 move.¹² If \mathcal{D} is not in memory, we must compute the fluence-normalized doses $D(v, \mathbf{b})$ for all $v \in \mathcal{V}(\mathbf{b}'_\ell)$ to obtain v^* , making the deposited doses $D(v, \mathbf{b}'_\ell) = n_{\omega_\ell}(\tau, e) D(v, \mathbf{b})$ simultaneously available to evaluate the objective variation (5.5). Hence the number of flops required in addition to $N_1(\omega, \omega') + N_{\mathcal{D},1}(\omega, \omega')$ is also $4|\mathcal{V}(\mathbf{b}'_\ell)|$, which is negligible compared to $N_{\mathcal{D},1}(\omega, \omega')$.

5.4 Type-3 moves

Unlike type-2 moves, type-3 moves are much more expensive than type-1 moves, because the optimal fluence must be computed for each energy considered. Let τ be the trajectory of a candidate beam \mathbf{b}'_ℓ defining a move from ω to $\omega' \in \mathcal{S}_3(\omega)$. Since the size of an energy partition block $\mathcal{E}_{p,j}(\tau)$ is $\lfloor p^{-1}|\mathcal{E}(\tau)| \rfloor$ or $\lceil p^{-1}|\mathcal{E}(\tau)| \rceil$, we can simply assume that $|\mathcal{E}_{p,j}(\tau)| = \lceil p^{-1}|\mathcal{E}(\tau)| \rceil$ for all j .

If \mathcal{D} is in memory, the cost of a type-3 move is that of a type-1 move plus that of computing the optimal fluences $n_{\omega_\ell}(\tau, e)$ for all $e \in \mathcal{E}_{p,j}(\tau)$; so the number of flops for implementing a type-3 move is

$$N_{3,p}(\omega, \omega') \approx N_1(\omega, \omega') + 4|\mathcal{V}(\mathbf{b}'_\ell)| \lceil p^{-1}|\mathcal{E}(\tau)| \rceil. \tag{5.20}$$

If \mathcal{D} is not in memory, we must compute the fluence-normalized doses $D(v, \mathbf{b})$, $\mathbf{b} := (\tau, e, 1)$, for all $v \in \mathcal{V}(\mathbf{b}'_\ell)$ and $e \in \mathcal{E}_{p,j}(\tau)$. To do this, we exploit the fact that the entry and exit depths $d_{\text{in}}[v, \mathbf{b}'_\ell(i)]$ and $d_{\text{out}}[v, \mathbf{b}'_\ell(i)]$ need to be computed only once regardless of the energy. The number of flops required in addition to $N_{3,p}(\omega, \omega')$ is then

$$N_{\mathcal{D},3,p}(\omega, \omega') \approx (3 \lceil p^{-1}|\mathcal{E}(\tau)| \rceil + N_d) \sum_{i \in \mathcal{I}} |\mathcal{V}(\mathbf{b}'_\ell(i))|. \tag{5.21}$$

5.5 Dynamic exploration strategy

We conclude the computational complexity analysis with the dynamic exploration strategy based on the communication sequence $(\Theta_{3,p_m,\eta})_m$ described in Sect. 4.4. The communication matrix $\Theta_{3,p,\eta}$ consists mainly of type-3 moves and incorporates type-1 moves with a small probability η ; so the average computational cost for implementing a move using $\Theta_{3,p,\eta}$ is almost the same as that of a type-3 move. The costs given below are average numbers of flops, as indicated by the overline notation.

¹² Most of the time, $\mathcal{V}(\mathbf{b}_\ell) \cap \mathcal{V}(\mathbf{b}'_\ell) = \emptyset$ or $|\mathcal{V}(\mathbf{b}_\ell) \cap \mathcal{V}(\mathbf{b}'_\ell)| \ll |\mathcal{V}(\mathbf{b}_\ell)| + |\mathcal{V}(\mathbf{b}'_\ell)|$ (unless the center axes of \mathbf{b}_ℓ and \mathbf{b}'_ℓ are parallel and close to each other), and so $N_1(\omega, \omega') \approx 4|\mathcal{V}(\mathbf{b}_\ell)| + 5|\mathcal{V}(\mathbf{b}'_\ell)|$.

Let $|\overline{\mathcal{V}(\mathbf{b})}|$ denote the average number of voxels intersected by a beam over all allowed trajectories. If \mathcal{D} is in memory, then, by (5.15), the cost of a type-1 move is $\overline{N}_1 \approx 9|\overline{\mathcal{V}(\mathbf{b})}|$. It follows from (5.20) that the cost of a type-3 move is

$$\overline{N}_{3,p} \approx (9 + 4p^{-1}|\overline{\mathcal{E}}|)|\overline{\mathcal{V}(\mathbf{b})}|. \tag{5.22}$$

Let R denote the beam-to-voxel cross-section ratio:

$$R := \pi(\zeta^{-1}\rho_{\max})^2, \tag{5.23}$$

where ζ is the voxel size and ρ_{\max} is the beam radius. If \mathcal{D} is not in memory, then, by (5.21), the additional cost is

$$\overline{N}_{\mathcal{D},3,p} \approx (3p^{-1}|\overline{\mathcal{E}}| + N_d)R^{-1}|\mathcal{I}||\overline{\mathcal{V}(\mathbf{b})}| \tag{5.24}$$

(we assume that the voxels are isotropic). Since $p^{-1}|\overline{\mathcal{E}}|$ ranges between 2 and $|\overline{\mathcal{E}}|$ during the optimization process, it follows that

$$\frac{3}{4}R^{-1}|\mathcal{I}| < \frac{\overline{N}_{\mathcal{D},3,p}}{\overline{N}_{3,p}} \leq \frac{6 + N_d}{17}R^{-1}|\mathcal{I}|. \tag{5.25}$$

So storing \mathcal{D} reduces the computational cost by a factor of the order of $R^{-1}|\mathcal{I}|$.

6 Tuning of the cooling schedule

Tuning a finite exponential cooling schedule $(T_{m,M})_{m \in [1..M]}$ [see (3.13)] consists in choosing the initial and final temperatures $T_{1,M} =: T_{\max}$ and $T_{M,M} =: T_{\min}$. This task is central to the performance of SA: if T_{\max} is too low, the behavior of the SA algorithm is close to deterministic, which usually leads to poor local minima; and if T_{\min} is too high, the final state distribution $\omega \in \Omega \mapsto P(X_M = \omega)$ is not concentrated enough around deep minima, meaning SA may terminate in a state not even close to a local minimum.

Practice shows that the probability to accept uphill moves should be close to one at the beginning of the annealing process (for efficiently exploring the state space at high temperatures) and close to zero when approaching the horizon M (for encouraging moves toward nearby local minima when freezing). These observations form the basis of our criteria for choosing T_{\max} and T_{\min} .

6.1 Criteria for choosing the initial and final temperatures

Recall that SA is defined by a family of transition matrices Q_T of the form (3.7). For any $T > 0$, we let $\mathcal{M}(Q_T)$ denote the homogeneous Markov chain with transition matrix Q_T [in other words, $\mathcal{M}(Q_T)$ is the Metropolis chain with landscape (Ω, Θ, U) and temperature T]. We define the *acceptance rate* $\zeta(T)$ of $\mathcal{M}(Q_T)$ to be the ratio of the probability of moving uphill to that of proposing a move uphill when $\mathcal{M}(Q_T)$ is at equilibrium; that is,

$$\zeta(T) := \frac{\sum_{(\omega, \omega') \in \Upsilon} \mathbf{v}_T(\omega) \Theta(\omega, \omega') \exp(-T^{-1}(U(\omega') - U(\omega)))}{\sum_{(\omega, \omega') \in \Upsilon} \mathbf{v}_T(\omega) \Theta(\omega, \omega')}, \tag{6.1}$$

where $\Upsilon := \{(\omega, \omega') \in \Omega^2 \mid U(\omega) < U(\omega')\}$ and \mathbf{v}_T is the stationary distribution of Q_T (this distribution is unique since Q_T inherits the irreducibility of the communication matrix Θ). Obviously, ζ is an increasing function of the temperature. Our approach to choosing

the initial and final temperatures is as follows: given target acceptance rates ζ_{\max} and ζ_{\min} such that $0 < \zeta_{\min} \ll \frac{1}{2} < \zeta_{\max} < 1$, find T_{\max} and T_{\min} such that $\zeta(T_{\max}) = \zeta_{\max}$ and $\zeta(T_{\min}) = \zeta_{\min}$.

It is worth noting that SA with exponential cooling is not significantly affected if T_{\max} is too high or if T_{\min} is too low, meaning it performs well as long as the orders of magnitude of the extremal temperatures are correct.¹³ The implications are twofold. First, we have some latitude in setting the target rates (from our experience, we can safely take $\zeta_{\max} \in [0.7, 0.9]$ and $\zeta_{\min} \in [10^{-4}, 10^{-3}]$). Second, T_{\max} and T_{\min} can be estimated using fast approximate methods such as those proposed in [59]. However, a fair comparison of the proposed communication strategies requires more accurate temperature estimates, which we now describe.

6.2 Practical tuning method

Empirical evidence suggests that the relationship between the acceptance rate and the temperature has the form

$$\zeta(T) \approx \begin{cases} 1 - a_1/T^{b_1} & \text{when } \zeta(T) \geq \frac{1}{2}, \\ a_0 T^{b_0} & \text{when } \zeta(T) \text{ is close to } 0, \end{cases} \tag{6.2}$$

where (a_1, b_1) and (a_0, b_0) are pairs of positive constants depending on the landscape. This model yields the initial and final temperature estimates

$$\hat{T}_{\max} := \left(\frac{a_1}{1 - \zeta_{\max}} \right)^{1/b_1} \quad \text{and} \quad \hat{T}_{\min} := \left(\frac{\zeta_{\min}}{a_0} \right)^{1/b_0}, \tag{6.3}$$

and now the problem is to determine (a_1, b_1) and (a_0, b_0) .

Suppose we can estimate the acceptance rates $\zeta_l := \zeta(T_l)$ at logarithmically spaced temperatures

$$T_l := 2^{p-l} T^*, \quad l \in [1 \dots q], \tag{6.4}$$

where the integers $p \geq 1$ and $q \geq 2$ and the temperature T^* are such that

$$\zeta_1 = \zeta(2^{p-1} T^*) \geq \zeta_{\max} \quad \text{and} \quad \zeta_q = \zeta(2^{p-q} T^*) \leq \zeta_{\min} \tag{6.5}$$

(a safe choice is to let T^* be a rough estimate of T_{\max} computed as described in [59], set $p = 4$, and adjust q accordingly). Then the pair (a_1, b_1) is easily obtained by fitting a regression line to the points $(\log T_l, \log(1 - \zeta_l))$ such that $\zeta_l \geq \frac{1}{2}$, and similarly for (a_0, b_0) with the points $(\log T_l, \log \zeta_l)$ such that $\zeta_l < \frac{1}{2}$.

Our approach to estimating ζ_1, \dots, ζ_q is as follows. Let $\omega_0 \in \Omega$ be picked at random, and let $(M_l)_{l \in \mathbb{N}}$ be a strictly increasing sequence in \mathbb{N}^* , with $M_0 = 1$. We monitor a realization $(\omega_m)_m$ of the SA algorithm with cooling schedule

$$(2^{p-l} T^*)_{m \in \mathbb{N}^*}, \quad l_m := 1 + \max \{l \in \mathbb{N} : M_l \leq m\} \tag{6.6}$$

(that is, starting from T_1 , the temperature is divided by 2 each time $m = M_l$ for some $l \geq 1$). For each $l \geq 1$, the acceptance rate ζ_l is estimated from the portion of $(\omega_m)_m$ at the l th temperature stage using the estimator

$$\hat{\zeta}_l := \frac{\sum_{m \in [M_{l-1} \dots M_l - 1]} \mathbb{1}_{\{U(\omega_{m-1}) < U(\omega_m)\}} \exp(-T_l^{-1}(U(\omega_m) - U(\omega_{m-1})))}{\sum_{m \in [M_{l-1} \dots M_l - 1]} \mathbb{1}_{\{U(\omega_{m-1}) < U(\omega_m)\}}}. \tag{6.7}$$

¹³ This observation does not apply to SA with logarithmic cooling.

The iterations are terminated when $\widehat{\zeta}_l \leq \zeta_{\min}$, determining the number of regression points q . We call this simulation process the *temperature-tuning chain*, so as to distinguish it from the tuned SA algorithm that subsequently minimizes the objective. The longer the lengths $M_l - M_{l-1}$ of the temperature stages, the sharper the estimates $\widehat{\zeta}_l$, and so the closer $\zeta(\widehat{T}_{\max})$ and $\zeta(\widehat{T}_{\min})$ to the target acceptance rates. Also, since the mixing time of $\mathcal{M}(Q_T)$ increases with decreasing temperature, the difference $M_l - M_{l-1}$ should increase with l . A reasonable choice is $M_l - M_{l-1} = \mu^{l-1}L$ with $\mu \in]1, 2]$ and $L \in \mathbb{N}^*$ such that the length of the temperature-tuning chain, $M_q = 1 + (\mu^q - 1)L/(\mu - 1)$, is some fraction of the horizon M . We recommend that the minimum length L of the temperature stages be of the order of the maximum communication distance between two states, so a sufficient number of uphill transitions are available to estimate the first regression point. We also found empirically that setting $\mu = 1.5$ is a good compromise between the accuracy of the acceptance rates and the computational cost of simulating the temperature-tuning chain. In fact, taking μ close to 1 provides effective cooling (meaning correct orders of magnitude for T_{\max} and T_{\min}), but not enough accuracy to properly compare the performance of different SA algorithms. On the other hand, increasing μ above 1.5 raises the accuracy to a level that has no measurable impact on performance.

Finally, we note that the temperature-tuning chain must be slightly modified when using a temperature-dependent exploration mechanism. In this case, the simulation process starts with the initial communication matrix (e.g., $\Theta_{3,p_{\max},\eta}$ for the exploration strategy described in Sect. 4.4) and switches to the final one (e.g., $\Theta_{3,1,\eta}$) as soon as $\widehat{\zeta}_l < \frac{1}{2}$.¹⁴

7 Numerical experiments

The experiments presented in this section demonstrate the feasibility of full ITP and highlight the benefits of the proposed exploration strategy over less sophisticated ones. We also show the benefits of full ITP over SIO and discuss the effect of the number of treatment beams per fraction on the delivered dose distribution.

It is important to note that performance is defined not only by the quality of the computed treatment plans (which can be measured using the objective function and other statistics of the delivered dose distributions), but also by the computational load. So the underlying goal is to find the best possible balance between quality and CPU time.

7.1 Irradiation geometries

We consider carbon-ion beams along with the irradiation geometries sketched in Fig. 2. The calculation volume is a set of cubic voxels of size 1 mm, with the PTV lying in a $2 \times 3 \times 4$ cm³ rectangular box \mathcal{B} centered in a cylindrical volume of water with a radius of 60 mm; the OAR region is partially surrounded by the PTV for increased difficulty. Each beam is modeled by $|\mathcal{I}| = 441$ infinitesimal rays (see Sect. 5.1) and has a 4 mm FWHM and a diameter of 8 mm (the fluence contributions of the peripheral rays are about 6% of that of the center ray). Furthermore, we simulate the use of a ripple filter so that the depth-dose profiles are approximately Gaussian in their peak region [65].

The first geometry in Fig. 2b, called the *reference geometry*, is specifically adapted to the shape of the PTV. It consists of $|\mathcal{T}| = 4.5 \times 10^4$ possible trajectories divided into four

¹⁴ The point $(\log T_l, \log \widehat{\zeta}_l)$ such that $\widehat{\zeta}_{l-1} \geq \frac{1}{2}$ and $\widehat{\zeta}_l < \frac{1}{2}$ must then be discarded from the regression data.

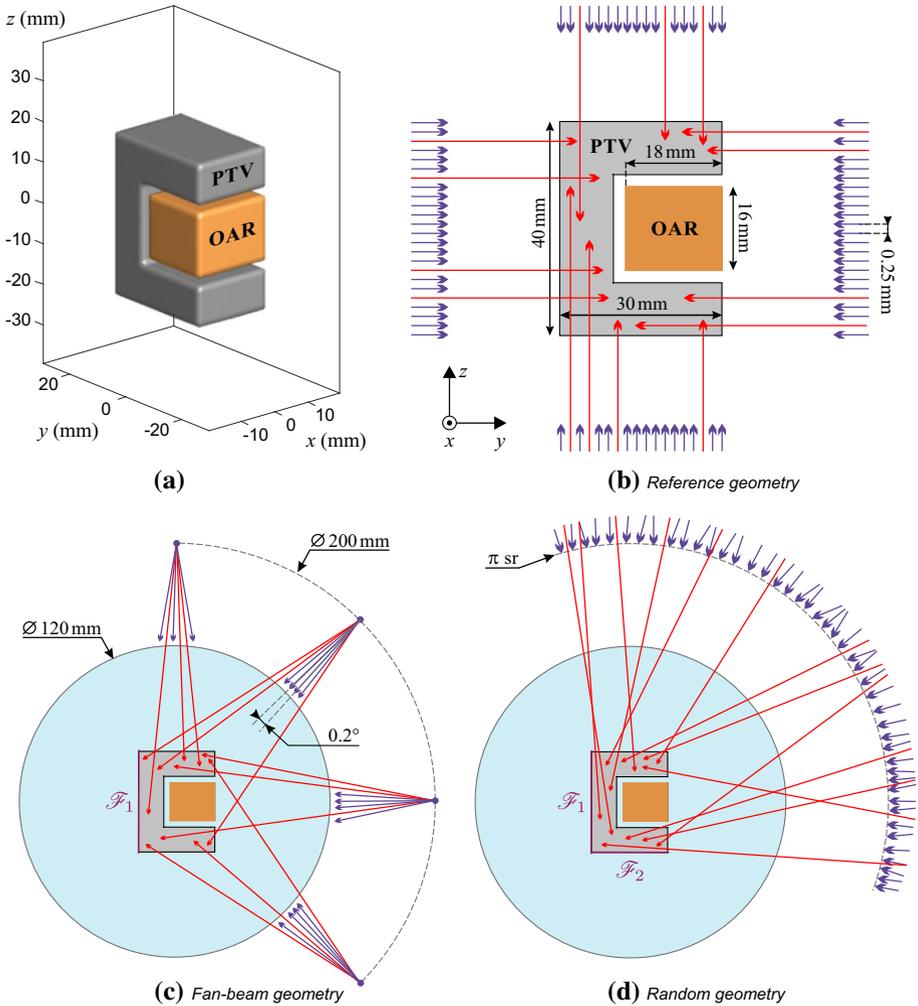


Fig. 2 Numerical phantom and irradiation geometries: **a** PTV and OAR region; **b** reference geometry adapted to the PTV (the small blue arrows schematize the allowed beam trajectories, of which only a small subset constitutes a treatment plan); **c** fan-beam geometry (the face \mathcal{F}_1 is not irradiated); **d** random geometry (the faces \mathcal{F}_1 and \mathcal{F}_2 are scarcely irradiated with small angles of incidence). (Color figure online)

subsets; in each subset, the trajectories are perpendicular to a face of \mathcal{B} , orthogonal to the x -axis, and evenly distributed with a spacing of 0.25 mm (along either the x - and y -directions or the x - and z -directions).

The second geometry in Fig. 2c, called the *fan-beam geometry*, is obtained by “stacking” identical 2-D fan-beam geometries along the x -axis with a step size of 0.25 mm. Each 2-D fan-beam geometry consists of four sources positioned on a circle of radius 100 mm with a 45° step angle. For each source, the beam angle increment is 0.2°, totaling $|\mathcal{T}| = 4.2 \times 10^4$ possible trajectories.

The third geometry in Fig. 2d, called the *random geometry*, consists of $|\mathcal{T}| = 400^2 = 1.6 \times 10^5$ trajectories generated from 400 source positions randomly sampled on a spherical

cap subtending a solid angle of π steradians (this cap is centered on \mathcal{B} and has a radius of 100 mm) and 400 points randomly sampled in the PTV.

The artificial shape of the PTV makes ITP particularly challenging for the fan-beam and random geometries, since obtaining a uniform dose distribution at the corners is difficult when the allowed beam trajectories are not perpendicular to the faces of the PTV. The fan-beam and random geometries have a further disadvantage in that they do not directly irradiate all three faces of the PTV opposite to the OAR. Therefore it will be no surprise that the best performance is obtained for the reference geometry. The reason for considering the fan-beam and random geometries is to show that full ITP can produce acceptable treatment plans in a nonstandard setting. In carbon-ion therapy, fan-beam geometries can be implemented in treatment facilities with a rotating gantry and fast magnetic deflection units (see, e.g., [1]). Random geometries are more prospective but will likely be implementable in the future thanks to the ongoing progress in accelerator technologies [67].

Finally, note that in all three geometries, a significant number of beam trajectories enter through the OAR. This is a deliberate choice to show that full ITP spares the OAR despite potentially harmful trajectories.

7.2 Annealing algorithms

The state space Ω [see (2.3)] is defined by an energy increment δ_e of 1 MeV u^{-1} (the corresponding depth resolution is about 0.625 mm), a fluence increment δ_n of 10^4 particles, and a maximum fluence n_{\max} of 5×10^6 particles. The dose plan D_\star in the ITP objective $U : \Omega \rightarrow \mathbb{R}$ [see (2.9)] is 2 Gy in the PTV and 0 in the OAR. Except for Sect. 7.3.4, the number of beams K is set to 600, and each SA algorithm specified below is run eight times to obtain an 8-fraction treatment plan for a target prescription dose of 16 Gy.¹⁵ In the case of the random geometry, each run uses a different set of possible trajectories, as we found that resampling \mathcal{T} for each fraction yields a smoother dose distribution in the PTV.

We consider the increasingly sophisticated SA algorithms listed in Table 2,¹⁶ so we can compare logarithmic versus exponential cooling and assess the efficiency of the communication mechanisms described in Sect. 4. Each algorithm starts from a random state $\omega_0 \in \Omega$ (i.e., a random feasible treatment plan). The probability η of using type-1 moves in $\mathbb{A}_{2,\text{exp}}$ and $\mathbb{A}_{3,\text{opt}}$ is set to 0.01, a choice that ensures irreducibility and symmetric support without compromising the benefits of type-2 and type-3 moves. In $\mathbb{A}_{3,\text{opt}}$, the parameter β controlling the decreasing rate of p_m (the number of energy partition blocks) is set to 0.2 to balance speed and accuracy. These choices of the parameters η and β are empirical, but there is some freedom in setting them. Indeed, as long as M is large enough (of the order of $10^3 K$, say) and the cooling schedule is properly tuned, we can arbitrarily choose $\eta \in]0, 0.1]$ and $\beta \in [0.1, 1]$ without significantly impacting performance (note that $\beta = 1$ is the threshold below which p_m decreases faster than linearly at each temperature change). Moreover, there is no sharp drop in performance when η and β are outside these ranges. For a fixed β , the performance decreases significantly only when $\eta > 0.5$, as $\mathbb{A}_{2,\text{exp}}$ and $\mathbb{A}_{3,\text{opt}}$ then behave more like $\mathbb{A}_{1,\text{exp}}$. Suppose now that η is fixed. In the limit $\beta \rightarrow 0$, we have $p_m = 1$ for all m , and so each type-3 move examines all possible energy values. On the other hand, in the limit $\beta \rightarrow +\infty$, we have $p_m = p_{\max}$ for all m , and so each type-3 move examines two energy

¹⁵ The total number of distinct beam trajectories in such a plan is therefore at most 4800, which remains much smaller than the number of possible trajectories $|\mathcal{T}|$.

¹⁶ The subscript “opt” in $\mathbb{A}_{3,\text{opt}}$ stands for “optimized” in reference to the incremental design of the proposed exploration strategy, and “p.c.” stands for “piecewise constant”

Table 2 Considered SA algorithms

Algorithm	Communication matrix		Cooling schedule
$\mathbb{A}_{1,\log}$	Θ_1	(Sect. 4.1)	Logarithmic (3.9)
$\mathbb{A}_{1,\exp}$	Θ_1		Exponential (3.12)
$\mathbb{A}_{2,\exp}$	$\Theta_{2,\eta}$	(Sect. 4.2)	Exponential
$\mathbb{A}_{3,\text{opt}}$	$\Theta_{3,p_m,\eta}$	(Sects. 4.3–4.4)	P.c. exponential (3.13)

values on average. In these two limit cases, $\mathbb{A}_{3,\text{opt}}$ outperforms $\mathbb{A}_{2,\exp}$, but less markedly than for $\beta \in [0.1, 1]$.

Note that since the neighborhood systems associated with the mixed communication mechanisms $\Theta_{2,\eta}$ and $\Theta_{3,p,\eta}$ are identical to that of Θ_1 ,¹⁷ the different SA algorithms have the same local minima with the same depths. Also, the communication distance is bounded by K .

Following the recommendations in Sect. 6, the initial and final temperatures of exponential cooling are selected to match the target acceptance rates $\zeta_{\max} = 0.8$ and $\zeta_{\min} = 5 \times 10^{-3}$, respectively, using a temperature-tuning chain with parameters $\mu = 1.5$ and $L = K$. In the piecewise-constant case, we change the temperature every $2K$ iterations, that is, we set the number of constant-temperature stages ξ to $\lceil M/(2K) \rceil$. (The choice of ξ , however, is not critical: from our experience, a value of the order of 100 or larger yields stable performance.) If the cooling is logarithmic, the final temperature is set to the same value as that of exponential cooling, and thus the initial temperature is fixed by the horizon M [since $T_{\max} = T_{\min} \log(M+1)/\log 2$]; in this way, $\mathbb{A}_{1,\log}$ and $\mathbb{A}_{1,\exp}$ have the same final transition matrix $Q_{T_{\min}}$ and so can be compared fairly.

7.3 Treatment planning optimization results

7.3.1 Comparison in terms of objective

Figure 3 plots the final objective value $U(\omega_M)$ versus CPU time when the horizon M of the SA algorithms is increased exponentially.¹⁸ The CPU times correspond to Matlab implementations run on a single core of an Intel Xeon processor E5-2697 v2 with the set of fluence-normalized doses \mathcal{D} [see (5.4)] stored in memory¹⁹ (this requires about 27 GB for the reference geometry, 19 GB for the fan-beam geometry, and 65 GB for the random geometry).

Overall, the performance in terms of objective increases with the sophistication of the SA algorithm, which validates our incremental approach to designing the communication mechanism. We make the following specific observations. First, exponential cooling speeds up the optimization process by a factor of at least 10 relative to logarithmic cooling: the solutions computed by $\mathbb{A}_{1,\exp}$ in 15 mn or so have an average objective value smaller than

¹⁷ Because $\mathcal{S}_{3,p}(\omega) \subset \mathcal{S}_2(\omega) \subset \mathcal{S}_1(\omega)$ for all $p \in \mathbb{N}^*$ and $\omega \in \Omega$.

¹⁸ Plotting $U(\omega_M)$ versus M is of little interest because the different communication mechanisms have different computational costs (see Sect. 5).

¹⁹ When \mathcal{D} is not stored in memory, the CPU time is multiplied by about $7 \times R^{-1}|\mathcal{Z}|$ for $\mathbb{A}_{1,\log}$ and $\mathbb{A}_{1,\exp}$, $4.8 \times R^{-1}|\mathcal{Z}|$ for $\mathbb{A}_{2,\exp}$, and $1.8 \times R^{-1}|\mathcal{Z}|$ for $\mathbb{A}_{3,\text{opt}}$, where $R = 16\pi$ is the beam-to-voxel cross-section ratio [see (5.23)].

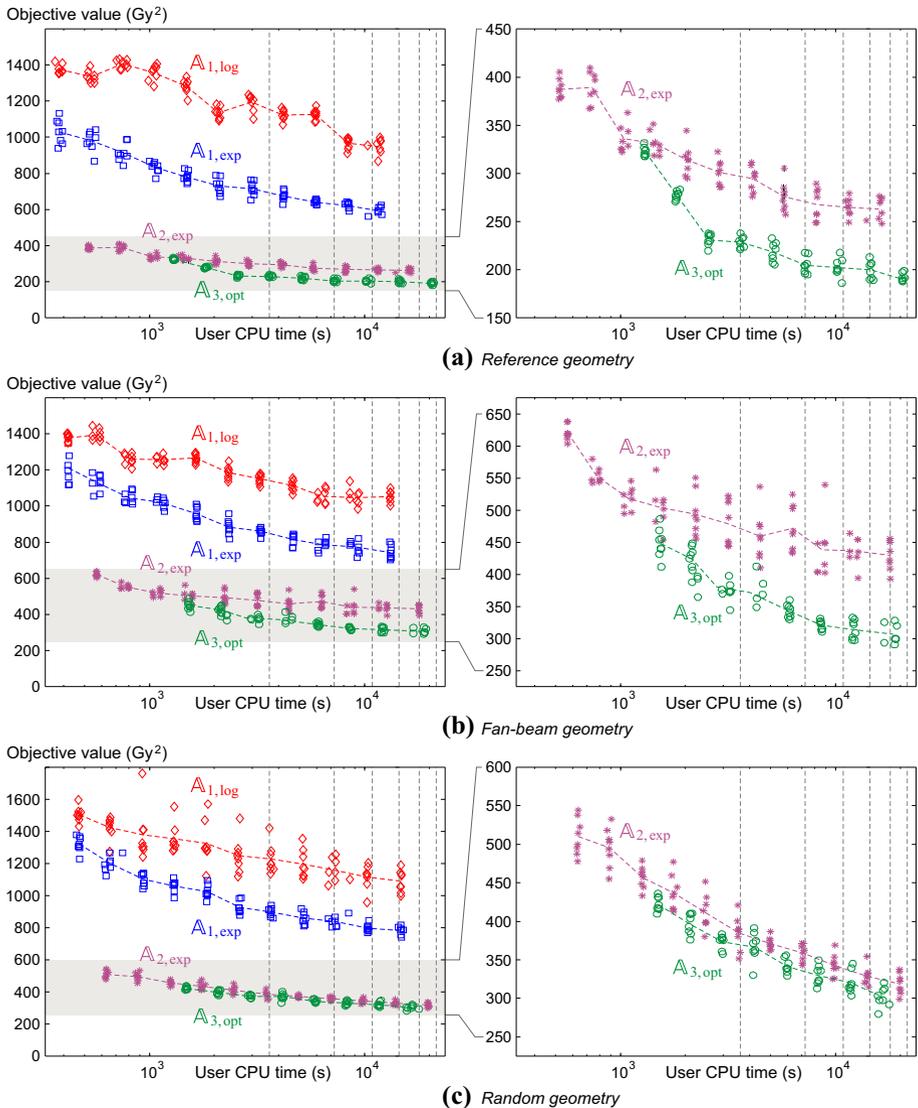


Fig. 3 Final objective value versus CPU time for the SA algorithms in Table 2: **a** reference geometry; **b** fan-beam geometry; **c** random geometry. Each algorithm is run eight times for each $M \in \{[2^{i/2} \times 100]; i \geq 5\}$ (the dashed curves connect the average objective values over each set of eight runs). The vertical dashed lines indicate one-hour periods for a single core of an Intel Xeon processor E5-2697 v2

that of the solutions computed by $\mathbb{A}_{1,\log}$ in about 3 h. Second, type-2 moves offer considerable performance improvements over type-1 moves: the objective values of the solutions computed by $\mathbb{A}_{2,\exp}$ in less than 12 mn are 15 to 30% lower than those of the solutions computed by $\mathbb{A}_{1,\exp}$ in 3 to 4 h, which corresponds to a speed-up factor greater than 16. In fact, extrapolating the average objective curves indicates that $\mathbb{A}_{2,\exp}$ is three orders of magnitude faster than $\mathbb{A}_{1,\exp}$. Third, the additional benefits brought by the optimized communication strategy used

in $\mathbb{A}_{3,\text{opt}}$ are less spectacular but still significant: for M sufficiently large, the speed-up factor over $\mathbb{A}_{2,\text{exp}}$ is greater than 8 for the reference and fan-beam geometries and about 1.5–2 for the random geometry. These speed-ups are achieved for M greater than about $10^3 K$ for the reference and fan-beam geometries and about $2.5 \times 10^3 K$ for the random geometry (these bounds correspond roughly to the third and fifth clusters of points in time order, respectively). Furthermore, as can be seen from the dispersion around the curves, $\mathbb{A}_{3,\text{opt}}$ is the most stable of the four algorithms in terms of objective variance.

The results in Fig. 3 also show that the SA landscapes have many poor local minima (namely, all the fraction plans computed by $\mathbb{A}_{1,\text{log}}$ and $\mathbb{A}_{1,\text{exp}}$). Furthermore, these local minima must be relatively deep, for otherwise there would not be such a performance gap between $\mathbb{A}_{1,\text{exp}}$ and $\mathbb{A}_{2,\text{exp}}$. From this viewpoint, the speed-ups achieved by $\mathbb{A}_{2,\text{exp}}$ and $\mathbb{A}_{3,\text{opt}}$ measure the increased ability of the mixed communication mechanisms to climb out of local basins.

7.3.2 Comparison in terms of dose distribution

Figures 4, 5 and 6 show, for each geometry, the dose distributions delivered by the 8-fraction treatment plans computed by $\mathbb{A}_{1,\text{log}}$, $\mathbb{A}_{1,\text{exp}}$ and $\mathbb{A}_{3,\text{opt}}$ in one hour or so. The dose distributions are represented as follows (from left to right): average doses along the coordinate axes; xy , yz and xz cross-sections; and 12 Gy isosurface.²⁰ We observe that exponential cooling and optimized communication both improve the dose homogeneity and yield sharper dose fall-off at the edges of the PTV. (The results obtained by $\mathbb{A}_{2,\text{exp}}$ are not shown because they are visually similar to those of $\mathbb{A}_{3,\text{opt}}$, but quantitative differences are highlighted below.) We also see that the OAR is almost completely spared by the treatment plans obtained using $\mathbb{A}_{3,\text{opt}}$; moreover, looking at the surroundings of the PTV shows that the trajectories that enter through the OAR are properly discarded.

Table 3 gives some statistics of the 8-fraction treatment plans computed by the different SA algorithms. These results should be interpreted in light of the fact that we use a smaller number of beams than in a standard clinical setting. Indeed, in SIO, the beam spacing in the directions perpendicular to the beam orientations and the average depth spacing should be about half the FWHM, that is 2 mm. For the 24 cm^3 PTV-OAR volume considered here, this requires 3×10^3 beams (or spots) per orientation, that is 9×10^3 beams for a typical number of 3 orientations—this is almost twice the number of beams of the presented plans (600 beams per fraction for a total of 4800 beams).

We see that increasing sophistication consistently improves performance in terms of distance to the dose plan and dose homogeneity: the minimum and maximum PTV dose get closer to the 16 Gy target prescription dose, the standard deviation of the PTV dose decreases, and so do the maximum, mean and upper quartile of the OAR dose.²¹ In agreement with the observations made in the previous section, the most important improvements come from the use of exponential cooling and type-2 moves, but the impact of the optimized communication strategy is also significant. Note that although the fraction plans produced by the different runs of a same SA algorithm have relatively close objective values (as seen in Fig. 3), they are far apart in terms of communication distance. A consequence is that the deviations of the

²⁰ An iso-value of 12 Gy allows to visualize the locations outside the PTV where the delivered dose is greater than 75% of the target prescription dose, which should be avoided as much as possible even though there is no constraint outside the PTV and the OAR.

²¹ Since the dose plan is zero in the OAR, the mean and upper quartile are better indicators of performance than the minimum and standard deviation.

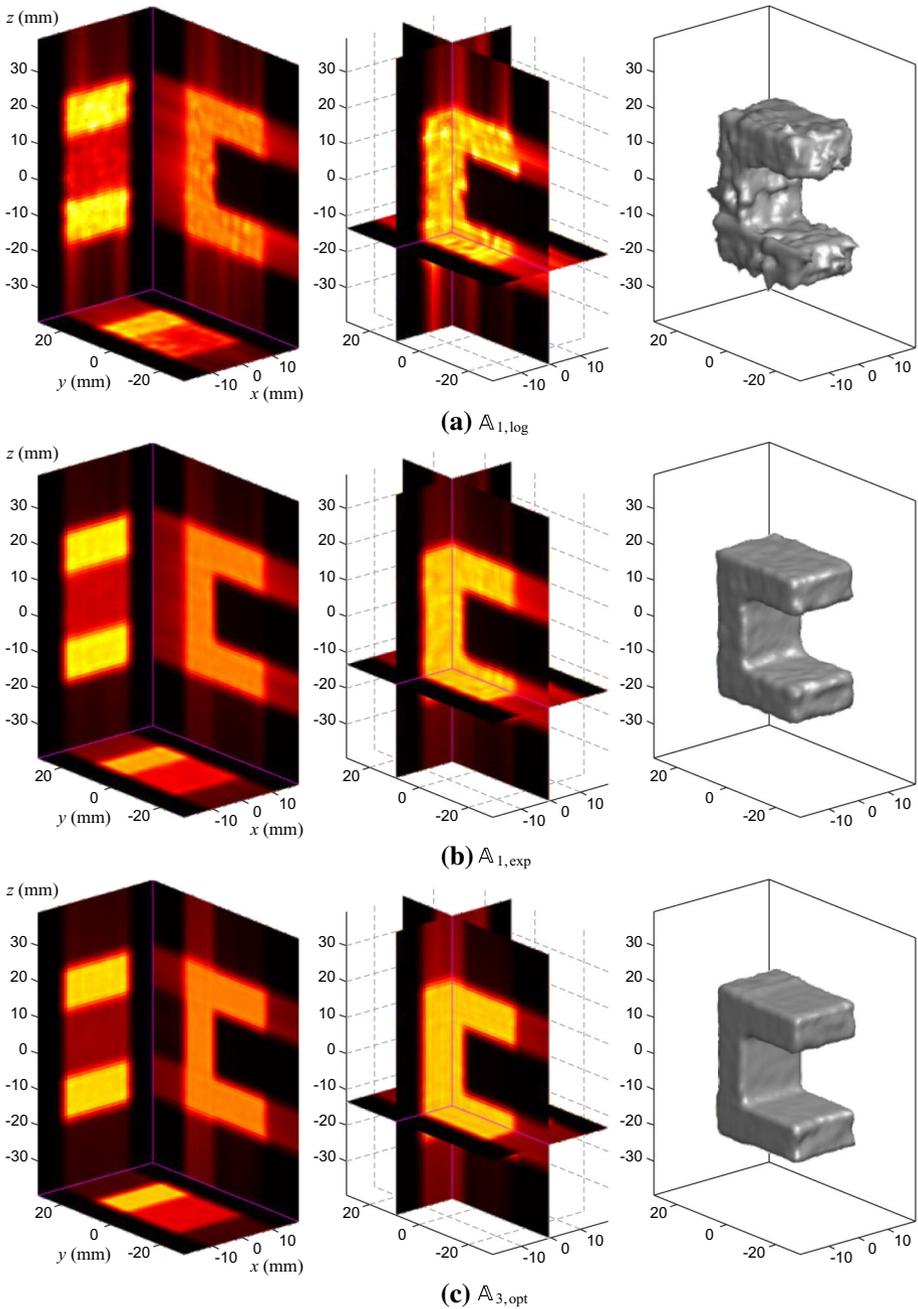


Fig. 4 Reference geometry: dose distributions delivered by the treatment plans computed by **a** $\mathbb{A}_{1,\log}$, **b** $\mathbb{A}_{1,\exp}$, and **c** $\mathbb{A}_{3,\text{opt}}$ ($M = 3200K$ for $\mathbb{A}_{3,\text{opt}}$ and $6400K$ otherwise). From left to right: average doses along the coordinate axes; xy , yz and xz cross-sections; and 12 Gy isosurface

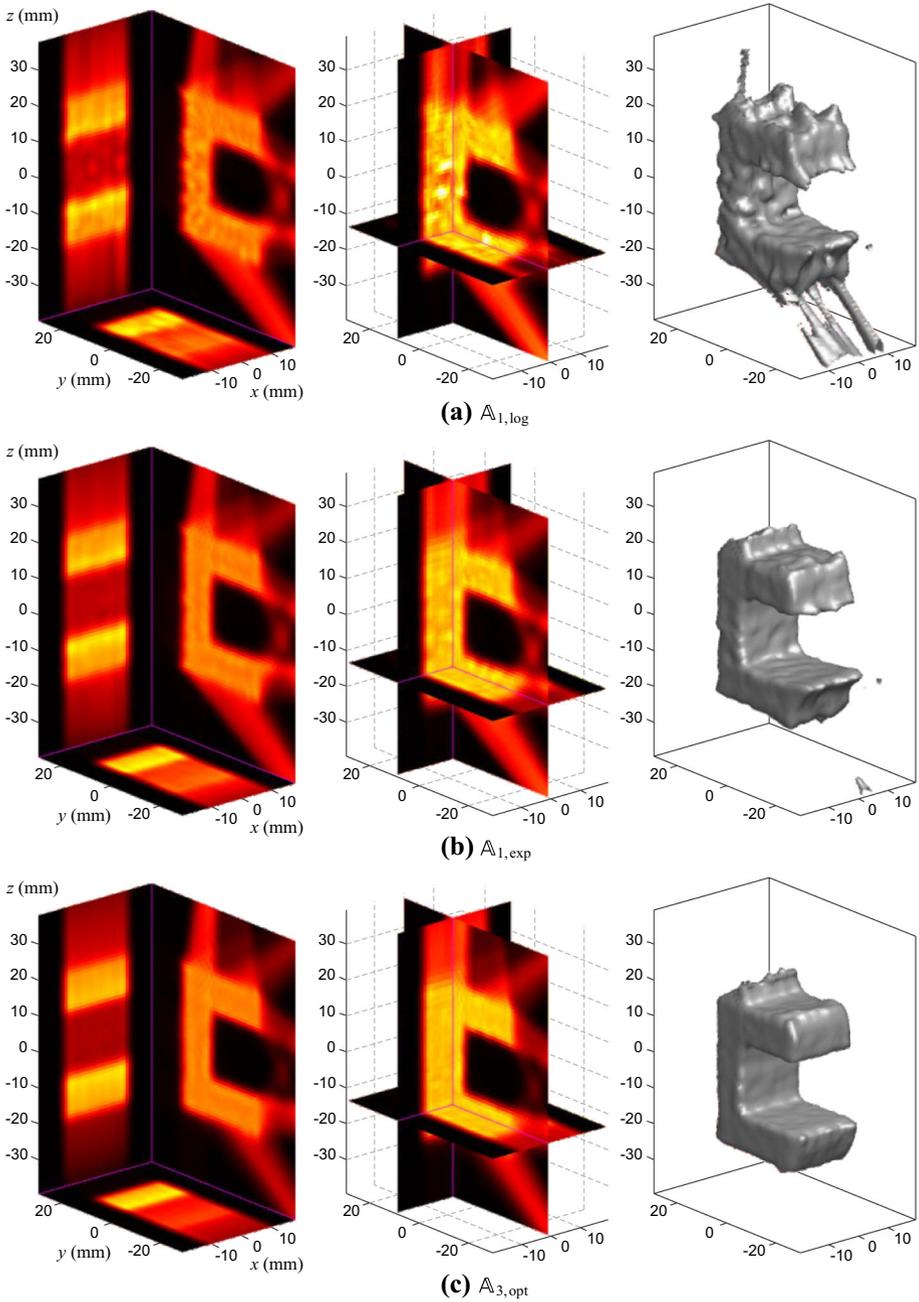


Fig. 5 Fan-beam geometry: dose distributions delivered by the treatment plans computed by **a** $\mathbb{A}_{1,\log}$, **b** $\mathbb{A}_{1,\exp}$, and **c** $\mathbb{A}_{3,\text{opt}}$ ($M = 3200K$ for $\mathbb{A}_{3,\text{opt}}$ and $6400K$ otherwise). From left to right: average doses along the coordinate axes; xy , yz and xz cross-sections; and 12Gy isosurface

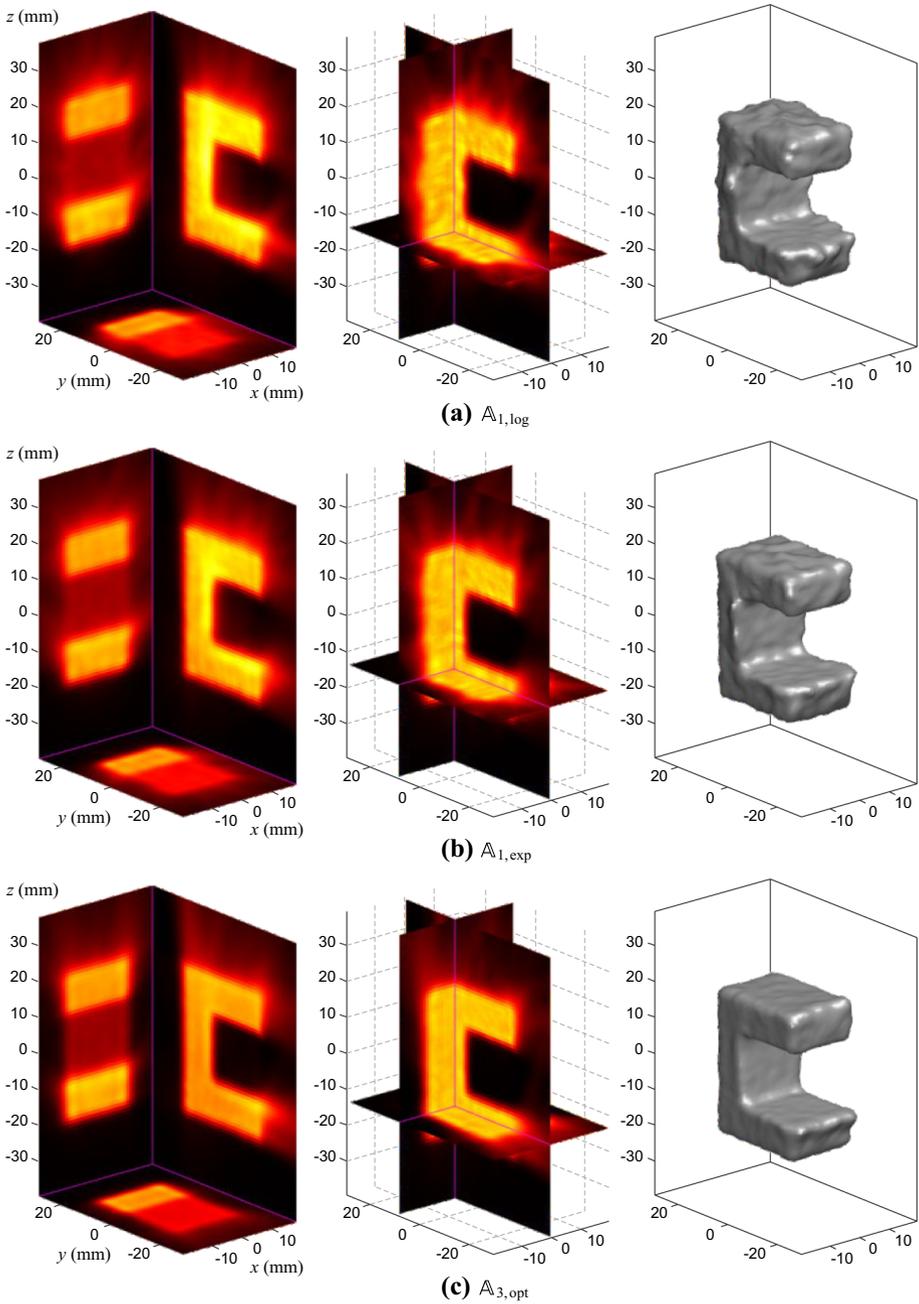


Fig. 6 Random geometry: dose distributions delivered by the treatment plans computed by **a** $\mathbb{A}_{1,\log}$, **b** $\mathbb{A}_{1,\exp}$, and **c** $\mathbb{A}_{3,\text{opt}}$ ($M = 3200K$ for $\mathbb{A}_{3,\text{opt}}$ and $6400K$ otherwise). From left to right: average doses along the coordinate axes; xy , yz and xz cross-sections; and 12Gy isosurface

Table 3 Statistics of the treatment plans computed by the different SA algorithms in similar amounts of time: average CPU time per fraction; root-mean-square distance (RMSD) to the dose plan; minimum, maximum and standard deviation of the PTV dose; maximum, mean and upper quartile (UQ) of the OAR dose

		Algorithm			
		$\mathbb{A}_{1,\log}$	$\mathbb{A}_{1,\text{exp}}$	$\mathbb{A}_{2,\text{exp}}$	$\mathbb{A}_{3,\text{opt}}$
Reference geometry					
PTV	CPU time	49 mn	50 mn	48 mn	43 mn
	RMSD	1.60	0.76	0.55	0.47
	Min.	7.74	12.78	13.17	14.10
	Max.	32.90	18.80	18.55	17.82
	Std. dev.	1.80	0.69	0.49	0.40
OAR	Max.	10.49	4.55	3.56	3.22
	Mean	1.00	0.83	0.50	0.45
	UQ	1.09	0.81	0.38	0.31
Fan-beam geometry					
PTV	CPU time	54 mn	55 mn	52 mn	52 mn
	RMSD	1.61	0.90	0.78	0.66
	Min.	8.57	11.75	12.04	13.04
	Max.	30.73	19.97	18.76	18.16
	Std. dev.	1.71	0.85	0.73	0.60
OAR	Max.	10.03	5.58	5.54	4.74
	Mean	0.98	0.77	0.59	0.56
	UQ	1.11	0.80	0.56	0.53
Random geometry					
PTV	CPU time	61 mn	61 mn	59 mn	49 mn
	RMSD	1.01	0.89	0.70	0.67
	Min.	10.95	11.93	12.70	12.75
	Max.	19.79	18.72	18.32	18.17
	Std. dev.	0.90	0.83	0.66	0.62
OAR	Max.	7.39	6.40	5.32	4.32
	Mean	0.93	0.80	0.53	0.53
	UQ	1.09	0.84	0.48	0.47

All dose statistics are expressed in Gy

fraction doses to the dose plan are poorly correlated, which has a smoothing effect on the total dose delivered. As an example, consider the eight fractions $\omega^{(1)}, \dots, \omega^{(8)}$ computed by $\mathbb{A}_{3,\text{opt}}$ for the reference geometry. A comparison of any two of these fractions shows that they have at most one identical beam, meaning they are at maximum or second-to-maximum communication distance from each other. The resulting smoothing effect is evidenced by the standard deviation of the total PTV dose, which falls from about 0.80 Gy when executing a same fraction $\omega^{(i)}$ eight times to 0.40 Gy when executing $\omega^{(1)}, \dots, \omega^{(8)}$.

Another way to assess the quality of a treatment plan is to look at the dose-volume histograms [68] of the OAR and the PTV. The dose-volume histogram of a region of interest \mathcal{R} (or \mathcal{R} -DVH, for short) is the complementary cumulative distribution of the dose in \mathcal{R} , that is, the function that assigns to each dose value D the relative volume of \mathcal{R} that receives

a dose greater than D . Ideally, the OAR-DVH should be 0 everywhere, and the PTV-DVH should be 1 (or 100%) if D is smaller than the prescribed dose and 0 otherwise. So the closer the OAR-DVH to the ordinate axis and the steeper the drop-off of the PTV-DVH at the prescribed dose, the better the treatment plan. Figure 7 shows, for each geometry, the OAR- and PTV-DVHs for the 8-fraction plans examined in Table 3. The difference in quality between two plans can be measured by the area between their respective DVHs. We observe that each added level of sophistication brings improvements in terms of both the OAR- and PTV-DVHs. (Note that the ripples in the DVHs associated with $\mathbb{A}_{2,\text{exp}}$ and $\mathbb{A}_{3,\text{opt}}$ for the reference geometry correspond to negligible modes of the dose distributions; indeed, since the negative derivative of the \mathcal{R} -DVH is the distribution of the dose in \mathcal{R} , a step-like variation in the DVH indicates a peak in the dose distribution.)

7.3.3 Full ITP versus SIO

In this section we show the benefits of full ITP over SIO. From now on, full ITP refers to computing a treatment plan using the algorithm $\mathbb{A}_{3,\text{opt}}$ described in Sect. 7.2.

The SIO geometry consists of three beam orientations with spherical angles $(90^\circ, 90^\circ)$, $(-90^\circ, 0^\circ)$ and $(90^\circ, -90^\circ)$. (These are the orientations entering the faces of the PTV opposite to the OAR in Fig. 2b, certainly the best choice for the PTV-OAR configuration considered.) The spots are distributed with a spacing of 2 mm in the directions perpendicular to each orientation and with a depth resolution of approximately 2.5 mm corresponding to a 4 MeV u^{-1} energy increment. The total number of beams—one for each spot—is 6780, which is about 40% more than for full ITP. The beam model is the same as before.

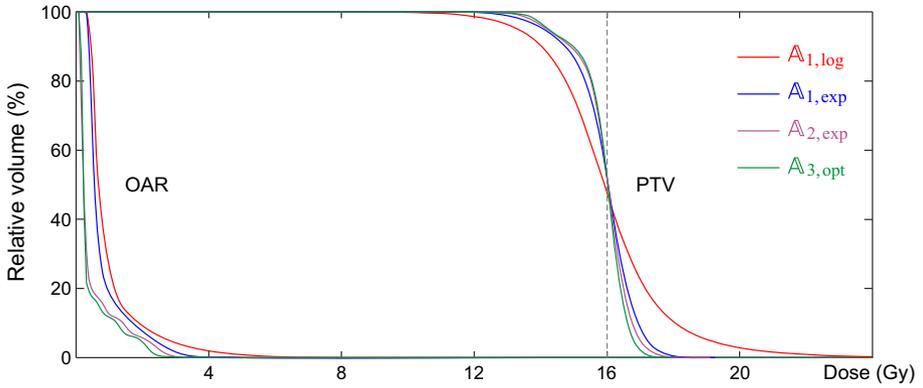
We performed SIO using SA with a communication mechanism that picks a spot u.a.r. and selects its candidate fluence value either u.a.r. or optimally (as in type-2 moves) with equal probability. (Note that SIO is easy for SA, because when the trajectories and energies are fixed, the objective is the restriction of a quadratic function to the finite set \mathcal{N}^K .) The dose distribution delivered by the resulting treatment plan is shown in Fig. 8; it is clearly not as smooth as those obtained by full ITP (compare with Figs. 4, 5 and 6c). Of course, the conformity to the dose plan improves with decreasing spot spacing; but we limit the number of spots for a fair comparison. The corresponding OAR- and PTV-DVHs are shown in Fig. 9 together with those obtained by full ITP for the reference geometry.

The statistics of the treatment plan computed by SIO are reported in Table 4. Comparing these results with those in the last column of Table 3, we see that, regardless of the geometry, full ITP performs much better than SIO in terms of all statistics except the maximum OAR dose (yet the difference is negligible in the case of the reference geometry). But that is not all. The last column of Table 4 gives the statistics of a 600-beam treatment plan obtained by full ITP for the reference geometry (we simply multiplied the fluence field of a single fraction $\omega^{(i)}$ by 8 to match the 16Gy target dose). Remarkably, this sparse treatment plan is superior to the one obtained by SIO, which shows that full ITP can drastically reduce the number of spots needed for tumor coverage.

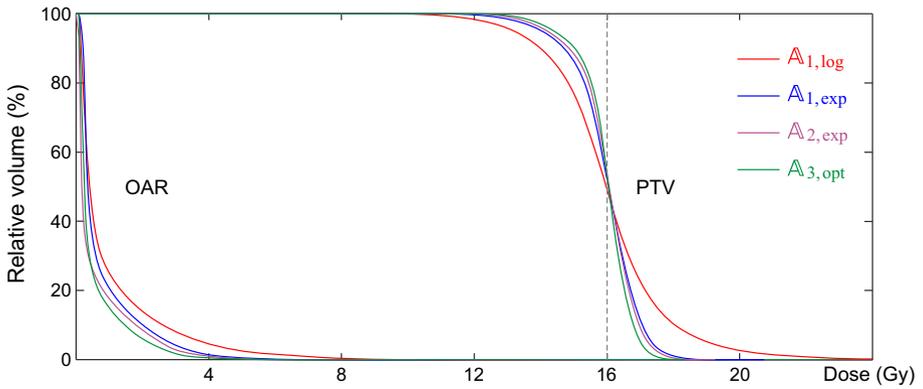
7.3.4 Effect of the number of beams

So far, the number of beams per fraction, K , has been kept constant, leaving aside the question of its influence on treatment plan quality. As we will see below, the conformity to the dose plan increases with K up to a point after which it remains approximately constant.

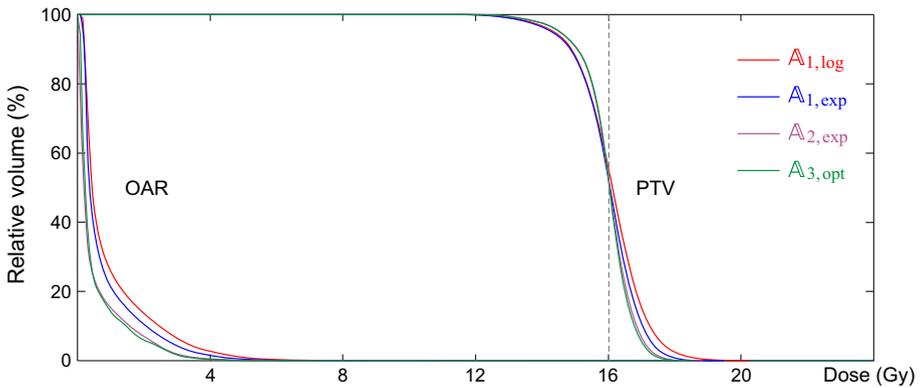
To assess the effect of the number of beams on dose homogeneity, we ran $\mathbb{A}_{3,\text{opt}}$ for $M = 3200K$ iterations with K ranging between 200 and 1400 (with a step of 25). Figure 10 shows



(a) Reference geometry



(b) Fan-beam geometry



(c) Random geometry

Fig. 7 OAR- and PTV-DVHs for the treatment plans computed by the different SA algorithms: **a** reference geometry; **b** fan-beam geometry; **c** random geometry. The vertical dashed lines indicate the target prescription dose

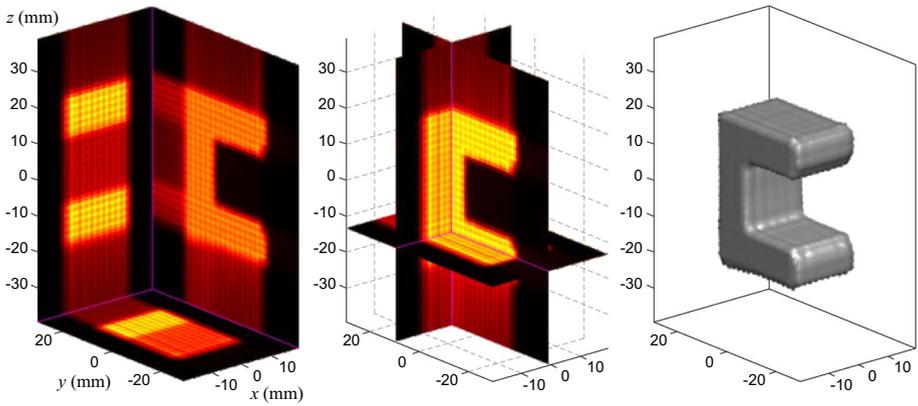


Fig. 8 Dose distribution delivered by the treatment plan obtained by SIO with 40% more beams than for full ITP (namely 6780 against 4800). From left to right: average doses along the coordinate axes; xy , yz and xz cross-sections; and 12 Gy isosurface

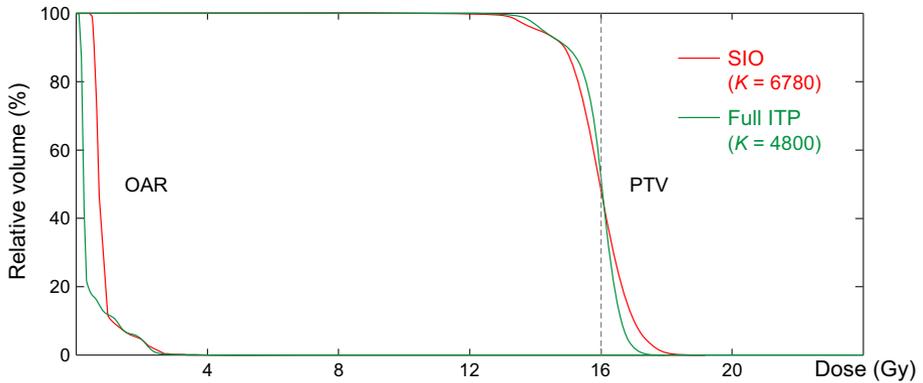


Fig. 9 OAR- and PTV-DVHs for the treatment plans obtained by SIO with 6780 beams and by full ITP for the reference geometry with 4800 beams

Table 4 Statistics of the 6780-beam treatment plan obtained by SIO and of the 600-beam treatment plan obtained by full ITP for the reference geometry: root-mean-square distance (RMSD) to the dose plan; minimum, maximum and standard deviation of the PTV dose; maximum, mean and upper quartile (UQ) of the OAR dose

		SIO ($K = 6780$)	Full ITP ($K = 600$)
PTV	RMSD	1.47	0.76
	Min.	9.45	12.08
	Max.	21.18	21.13
	Std. dev.	1.62	0.79
OAR	Max.	3.20	3.55
	Mean	0.81	0.44
	UQ	0.87	0.37

All dose statistics are expressed in Gy

the final objective value $U(\omega_M)$ and the standard deviation of the dose delivered by ω_M to the PTV as functions of K for the reference geometry (the results for the other geometries are

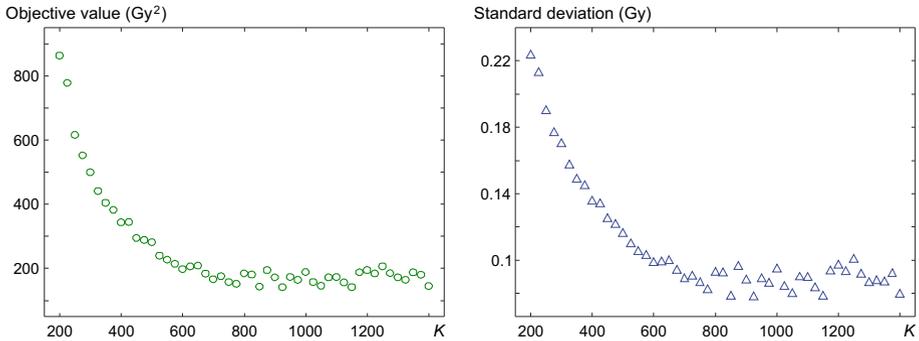


Fig. 10 Effect of the number of beams per fraction, K , on the final objective value and on the standard deviation of the PTV dose for the reference geometry. Each point corresponds to a single-fraction plan computed by $\mathbb{A}_{3,\text{opt}}$ with $M = 3200K$

Table 5 Statistics of the N_f -fraction treatment plans computed by $\mathbb{A}_{3,\text{opt}}$ (with $M = 3200K$) for the reference geometry and a constant total number of beams ($N_f K = 4800$): root-mean-square distance (RMSD) to the dose plan; minimum, maximum and standard deviation of the PTV dose; maximum, mean and upper quartile (UQ) of the OAR dose

	N_f	12	8	6	4
	K	400	600	800	1200
PTV	RMSD	0.52	0.47	0.46	0.61
	Min.	13.88	14.10	14.07	13.24
	Max.	17.97	17.82	21.72	25.21
	Std. dev.	0.47	0.40	0.39	0.61
OAR	Max.	3.17	3.22	3.09	3.07
	Mean	0.45	0.44	0.48	0.51
	UQ	0.33	0.32	0.36	0.43

All dose statistics are expressed in Gy and the best results are set in boldface

similar). These two quality measures are highly correlated because the objective is the square deviation between the delivered and prescribed dose, which is approximately proportional to the variance of the dose delivered to the PTV (provided, as is the case here, that the OAR receives negligible dose). We see that the quality of the computed treatment plan ω_M decreases with K and plateaus for $K \geq 800$. This shows that a horizon M linear in K does not compensate for the increase in optimization difficulty caused by increasing K above some threshold. Unfortunately, this threshold is geometry-dependent and costly to estimate, and so is the relationship between M and K yielding the best performance. As a rule of thumb, our general experience with annealing suggests that the horizon should be of the order of 10^3 times the number of variables in the objective. Therefore, since U has $3K$ variables (three parameters per beam), we recommend setting M between $10^3 K$ and $10^4 K$.

Let N_f denote the number of computed fractions. The existence of a performance limit when K increases and M is linear in K indicates that there is a computational compromise between N_f and K when the total number of treatment beams (i.e., $N_f K$) is fixed and the cumulative CPU time (i.e., the sum of the CPU times for all fractions) is bounded. For example, Table 5 shows the effect of N_f on the delivered dose when $N_f K = 4800$ and the fractions are computed by $\mathbb{A}_{3,\text{opt}}$ with $M = 3200K$ (so the cumulative CPU time is approximately constant). We see a significant decline in performance in terms of the RMSD, the PTV statistics and the upper quartile of the OAR dose when K increases from 800 to 1200, which agrees with the plateau observed in Fig. 10.

Finally, let us be clear that we do not suggest to choose the number of fractions based solely on optimization results — in practice, the fractionation scheme and the dose per fraction are prescribed by the radiation oncologist using insights from clinical trials and accounting for radiobiological effectiveness. Note, however, that a fractionation scheme optimized for best dose distribution can be easily adapted to clinical use: given a treatment plan \mathcal{T} with N_f different fractions computed by SA, we can define a new plan with N'_f identical fractions consisting of the union of all the beams of \mathcal{T} with their fluences scaled by $1/N'_f$. This new plan is uniform across fractions and leverages the smoothing effect described in Sect. 7.3.2.

8 Concluding remarks

We provided a proof of concept for full ITP in charged-particle therapy, which is to find optimal beam trajectories, energies and fluences to irradiate a target tumor while sparing nearby healthy structures. This combinatorial optimization problem is particularly challenging and, except for our work in [57], has not been previously addressed. We proposed an original SA approach whose exploration strategy exploits the linear dose-fluence relationship together with interlaced energy partitioning. We also introduced a new technique to accurately estimate the initial and final temperatures of the cooling schedule, which are crucial to the success of annealing.

We tested the performance of increasingly sophisticated SA algorithms for different irradiation configurations: a reference geometry specifically adapted to the shape of the PTV, and two generic geometries, the fan-beam and random geometries. The proposed approach consistently and substantially outperforms SA with uniform candidate beams as well as SA with optimal fluence moves. Moreover, the results obtained for the fan-beam and random geometries are fairly close to those for the reference geometry: the differences in the RMSD to the dose plan and in the standard deviation of the PTV dose do not exceed 0.22 Gy — less than 1.4% of the 16 Gy target dose — which is quite small in view of the fact that the fan-beam and random geometries are not tailored to the PTV and do not irradiate all three faces of the PTV opposite to the OAR. This suggests some robustness with respect to the choice of the irradiation geometry. We also observed significant improvements over SIO indicating the potential to greatly reduce the number of treatment beams.

We saw in Sect. 5.5 that the average computation time per iteration increases with $|\overline{\mathcal{E}}|$ and $|\overline{\mathcal{V}(\mathbf{b})}|$ (the average numbers of energies and beam-voxel intersections per trajectory), both of which increase approximately linearly with $|\mathcal{V}|^{1/3}$ (assuming fixed-size isotropic voxels). Using (5.22), it can be shown that for a fixed horizon M , the (total) computation time is approximately linear in $|\overline{\mathcal{V}(\mathbf{b})}|$ and affine in $\log |\overline{\mathcal{E}}|$, and therefore $O(|\mathcal{V}|^{1/3} \log |\mathcal{V}|)$. As discussed in Sect. 7.3.4, the horizon M must be at least linear in the number of treatment beams K ; furthermore, K should increase linearly with $|\mathcal{V}|$ for proper tumor coverage. So the computation time is at best $O(|\mathcal{V}|^{4/3} \log |\mathcal{V}|)$ and is currently of the order of one hour for small tumors (4–5 cm in diameter) and one day for large ones (8–10 cm in diameter). This is an obstacle for routine management of treatment planning. Potential remedies for achieving clinically acceptable computation times (i.e., 15–20 mn) include voxel clustering [69], voxel sampling [70], parallel SA by multiple trials [71], and parallel SA using speculative computation [72,73]. Moreover, the generation of type-3 candidate beams can be accelerated by parallel computation of the optimal fluences. Another limitation is the amount of memory required to store the set of fluence-normalized doses \mathcal{D} , that is approximately $\frac{4}{3} R |\mathcal{V}|^{1/3} |\mathcal{T}| |\overline{\mathcal{E}}|$ double precision elements. Storing \mathcal{D} reduces the computation time con-

siderably, but the downside is that the maximum number of allowed beam trajectories — and hence the flexibility of the irradiation geometry — is limited by the size of the tumor and the available memory.²²

Our framework can be extended to objective functions of the more general form

$$U(\omega) := \sum_{v \in \mathcal{V}} \left(\varphi_v(D_*(v)) - \sum_{k \in [1..K]} \varphi_{v, \tau_k, e_k}(D(v, \mathbf{b}_k)) \right)^2, \tag{8.1}$$

where $\{\varphi_v\}_{v \in \mathcal{V}}$ and $\{\varphi_{v, \tau, e}\}_{(v, \tau, e) \in \mathcal{V} \times (\mathcal{T} \otimes \mathcal{E})}$ are families of differentiable nondecreasing functions and $D(v, \mathbf{b})$ is defined as in Sect. 5.1. This allows to optimize the biological dose using the linear-quadratic model or the microdosimetric kinetic model (see, e.g., [74]). And if the dose model is computationally too expensive for stochastic optimization, the proposed SA algorithm can always be used with a simplified model for properly initializing a deterministic optimization pass with the full model.

Another avenue would be to augment the objective in order to enforce bounds on the PTV dose and to better control OAR sparing. Let $\mathcal{V}_0, \dots, \mathcal{V}_p$ be disjoint subsets of voxels, with \mathcal{V}_0 representing the PTV and $\mathcal{V}_1, \dots, \mathcal{V}_p$ representing different OARs. Let D_0 denote the prescribed dose and let $D(v) := \sum_{k \in [1..K]} D(v, \mathbf{b}_k)$ be the dose deposited in a voxel v by the k th beam of a treatment plan $\omega := (\mathbf{b}_1, \dots, \mathbf{b}_K)$. Define the augmented objective function

$$\begin{aligned} U^\dagger(\omega) := & \sum_{v \in \mathcal{V}_0} (D_0 - D(v))^2 + \lambda_0 \sum_{v \in \mathcal{V}_0} \max(D(v) - D_{\max}, D_{\min} - D(v)) \\ & + \sum_{j \in [1..p]} \lambda_j (c_j \max_{v \in \mathcal{V}_j} D(v) + (1 - c_j) \text{mean}_{v \in \mathcal{V}_j} D(v)), \end{aligned} \tag{8.2}$$

where the weights λ_j reflect the relative importance of the different structures, $D_{\min} < D_0$ and $D_{\max} > D_0$ are the prescribed minimum and maximum PTV doses, and each weight $c_j \in [0, 1]$ balances the maximum and mean doses in the j th OAR. Each term in the last sum corresponds to the linearized equivalent uniform dose model proposed in [75]. This model allows to distinguish between parallel and serial OARs by setting c_j close to 0 and 1, respectively.²³ To minimize U^\dagger , the main modification of our methodology would be in the computation of the optimal fluence for type-2 and type-3 candidate beams. We would suggest to search for optimal fluences with respect to the quadratic part of U^\dagger , that is,

$$Q^\dagger(\omega) := \sum_{v \in \mathcal{V}_0} (D_0 - D(v))^2 + \sum_{j \in [1..p]} \lambda_j (1 - c_j) \text{mean}_{v \in \mathcal{V}_j} D(v). \tag{8.3}$$

The reason is twofold. First, this strategy would not increase the computational cost. Second, since Q^\dagger is obtained from U^\dagger by removing the minimum and maximum dose constraints, Q^\dagger -optimal fluence moves leading to an excessive violation of these constraints would most likely be rejected in the mid- and low-temperature regimes, and hence inconsequential.

²² To give an idea, consider for simplicity a spherical tumor of radius r (mm), 1 mm^3 isotropic voxels, and a depth resolution of 1 mm. In this case \mathcal{D} occupies about $30 (r \rho_{\max})^2 |\mathcal{T}|$ bytes of memory, where ρ_{\max} is the beam radius. So for a 60 mm diameter tumor, a 4 mm beam radius, and 120 GB allocated to \mathcal{D} , the irradiation geometry is limited to approximately 1.5×10^5 trajectories.

²³ Parallel OARs tolerate high doses in small regions provided the mean dose received is small, whereas high doses are harmful to serial OARs even when confined to small regions. Examples of parallel OARs include the lung, kidney and liver; examples of serial OARs include the brain stem, spinal cord and esophagus.

As a final remark, we note that

$$\Omega \subset \Omega' \implies \min_{\omega \in \Omega} U(\omega) \geq \min_{\omega \in \Omega'} U(\omega) \quad (8.4)$$

for any sets Ω and Ω' of feasible treatment plans. Therefore, since $U(\omega)$ is the square deviation between the dose plan and the dose delivered by ω , increasing the number of allowed beam trajectories improves the quality of the treatment plans that achieve the global minima of the objective. This suggests gathering various irradiation geometries ($\mathcal{T}_1, \dots, \mathcal{T}_q$, say) and searching for optimal plans over $(\bigcup_{j \in [1..q]} (\mathcal{T}_j \otimes \mathcal{E}) \times \mathcal{N})^K$. However, the global geometry $\bigcup_{j \in [1..q]} \mathcal{T}_j$ is constrained by computational and memory resources and by the fact that current ion-beam delivery systems cannot execute all feasible plans (for example, those obtained from random geometries). Yet our numerical experiments show that benefits can be expected from introducing randomness into the irradiation geometry, which motivates research toward more flexible dose delivery systems in charged-particle therapy.

Acknowledgements This work was supported by the Chinese Basic Research Program under Grant 61671049 and the Chinese Key R&D Program under Grant 2017YFB1400100. We would like to thank Dr. François Smekens for providing us with the cumulative depth-dose profiles of carbon ions in water.

References

1. Schardt, D., Elsässer, T., Schulz-Ertner, D.: Heavy-ion tumor therapy: physical and radiobiological benefits. *Rev. Mod. Phys.* **82**(1), 383–425 (2010)
2. Kraft, G., Weber, U.: Tumor therapy with ion beams, pp. 1179–1205. Springer, Handbook of particle detection and imaging (2012)
3. Loeffler, J., Durante, M.: Charged particle therapy—optimization, challenges and future directions. *Nat. Rev. Clin. Oncol.* **10**(7), 411–424 (2013)
4. Schlaff, C., Krauze, A., Belard, A., O’Connell, J., Camphausen, K.: Bringing the heavy: carbon ion therapy in the radiobiological and clinical context. *Radiat. Oncol.* **9**, 88 (2014)
5. Durante, M., Orecchia, R., Loeffler, J.: Charged-particle therapy in cancer: clinical uses and future perspectives. *Nat. Rev. Clin. Oncol.* **14**, 483–495 (2017)
6. Shepard, D., Ferris, M., Olivera, G., Rockwell Mackie, T.: Optimizing the delivery of radiation therapy to cancer patients. *SIAM Rev.* **41**(4), 721–744 (1999)
7. Romeijn, H., Ahuja, R., Dempsey, J., Kumar, A., Li, J.: A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Phys. Med. Biol.* **48**(21), 3521–3542 (2003)
8. Romeijn, H., Ahuja, R., Dempsey, J., Kumar, A.: A new linear programming approach to radiation therapy treatment planning problems. *Oper. Res.* **54**(2), 201–216 (2006)
9. Chen, W., Craft, D., Madden, T., Zhang, K., Kooy, H., Herman, G.: A fast optimization algorithm for multicriteria intensity modulated proton therapy planning. *Med. Phys.* **37**(9), 4938–4945 (2010)
10. Cao, W., Lim, G., Li, X., Li, Y., Zhu, X., Zhang, X.: Incorporating deliverable monitor unit constraints into spot intensity optimization in intensity-modulated proton therapy treatment planning. *Phys. Med. Biol.* **58**(15), 5113–5125 (2013)
11. Lomax, A.: Intensity modulation methods for proton radiotherapy. *Phys. Med. Biol.* **44**(1), 185–205 (1999)
12. Krämer, M., Jäkel, O., Haberer, T., Kraft, G., Schardt, D., Weber, U.: Treatment planning for heavy-ion radiotherapy: physical beam model and dose optimization. *Phys. Med. Biol.* **45**(11), 3299–3317 (2000)
13. Oelfke, U., Bortfeld, T.: Inverse planning for photon and proton beams. *Med. Dosim.* **26**(2), 113–124 (2001)
14. Lahanas, M., Schreiber, E., Baltas, D.: Multiobjective inverse planning for intensity modulated radiotherapy with constraint-free gradient-based optimization algorithms. *Phys. Med. Biol.* **48**(17), 2843–2871 (2003)
15. Bourhaleb, F., Marchetto, F., Attili, A., Pittà, G., Cirio, R., Donetti, M., Giordanengo, S., Givehchi, N., Iliescu, S., Krenkli, M., La Rosa, A., Massai, D., Pecka, A., Pardo, J., Peroni, C.: A treatment planning

- code for inverse planning and 3D optimization in hadrontherapy. *Comput. Biol. Med.* **38**(9), 990–999 (2008)
16. Gemmel, A., Hasch, B., Ellerbrock, M., Weyrather, W., Krämer, M.: Biological dose optimization with multiple ion fields. *Phys. Med. Biol.* **53**(23), 6691–7012 (2008)
 17. Unkelbach, J., Bortfeld, T., Martin, B., Soukup, M.: Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning. *Med. Phys.* **36**(1), 149–163 (2009)
 18. Li, Y., Zhang, X., Mohan, R.: An efficient dose calculation strategy for intensity modulated proton therapy. *Phys. Med. Biol.* **56**(4), N71–N84 (2011)
 19. Inaniwa, T., Kanematsu, N., Furukawa, T., Noda, K.: Optimization algorithm for overlapping-field plans of scanned ion beam therapy with reduced sensitivity to range and setup uncertainties. *Phys. Med. Biol.* **56**(6), 1653–1669 (2011)
 20. Pflugfelder, D., Wilkens, J., Nill, S., Oelfke, U.: A comparison of three optimization algorithms for intensity modulated radiation therapy. *Z. Med. Phys.* **18**(2), 111–119 (2008)
 21. Liu, W., Zhang, X., Li, Y., Mohan, R.: Robust optimization of intensity modulated proton therapy. *Med. Phys.* **39**(2), 1079–1091 (2012)
 22. Romeijn, H., Ahuja, R., Dempsey, J., Kumar, A.: A column generation approach to radiation therapy treatment planning using aperture modulation. *SIAM J. Optim.* **15**(3), 838–862 (2005)
 23. Aleman, D., Glaser, D., Romeijn, H., Dempsey, J.: Interior point algorithms: guaranteed optimality for fluence map optimization in IMRT. *Phys. Med. Biol.* **55**(18), 5467–5482 (2010)
 24. Preciado-Walters, F., Langer, M., Rardin, R., Thai, V.: Column generation for IMRT cancer therapy optimization with implementable segments. *Ann. Oper. Res.* **148**, 65–79 (2006)
 25. Fredriksson, A.: A characterization of robust radiation therapy treatment planning methods—from expected value to worst case optimization. *Med. Phys.* **39**(8), 5169–5181 (2012)
 26. Hartmann, L., Bogner, L.: Investigation of intensity-modulated radiotherapy optimization with gEUD-based objectives by means of simulated annealing. *Med. Phys.* **35**(5), 2041–2049 (2008)
 27. Cotrutz, C., Xing, L.: Segment-based dose optimization using a genetic algorithm. *Phys. Med. Biol.* **48**(18), 2987–2998 (2003)
 28. Li, Y., Yao, J., Yao, D.: Genetic algorithm based deliverable segments optimization for static intensity-modulated radiotherapy. *Phys. Med. Biol.* **48**(20), 3353–3374 (2003)
 29. Ahmad, S., Bergen, S.: A genetic algorithm approach to the inverse problem of treatment planning for intensity-modulated radiotherapy. *Biomed. Signal Process. Control* **5**(3), 189–195 (2010)
 30. Cao, W., Lim, G., Liao, L., Li, Y., Jiang, S., Li, X., Li, H., Suzuki, K., Zhu, X., Gomez, D., Zhang, X.: Proton energy optimization and reduction for intensity-modulated proton therapy. *Phys. Med. Biol.* **59**(21), 6341–6354 (2014)
 31. Cao, W., Lim, G., Lee, A., Li, Y., Liu, W., Zhu, X., Zhang, X.: Uncertainty incorporated beam angle optimization for IMPT treatment planning. *Med. Phys.* **39**(8), 5248–5256 (2012)
 32. Lim, G., Kardar, L., Cao, W.: A hybrid framework for optimizing beam angles in radiation therapy planning. *Ann. Oper. Res.* **217**, 357–383 (2014)
 33. Gu, W., O'Connor, D., Nguyen, D., Yu, V., Ruan, D., Dong, L., Sheng, K.: Integrated beam orientation and scanning-spot optimization in intensity-modulated proton therapy for brain and unilateral head and neck tumors. *Med. Phys.* **45**(4), 1338–1350 (2018)
 34. Ehrgott, M., Güler, Ç., Hamacher, H., Shao, L.: Mathematical optimization in intensity modulated radiation therapy. *Ann. Oper. Res.* **175**, 309–365 (2010)
 35. Craft, D.: Local beam angle optimization with linear programming and gradient search. *Phys. Med. Biol.* **52**(7), N127–N135 (2007)
 36. Aleman, D., Romeijn, H., Dempsey, J.: A response surface approach to beam orientation optimization in intensity-modulated radiation therapy treatment planning. *INFORMS J. Comput.* **21**(1), 62–76 (2009)
 37. Zhang, H., Shi, L., Meyer, R., Nazareth, D., D'Souza, W.: Solving beam-angle selection and dose optimization simultaneously via high-throughput computing. *INFORMS J. Comput.* **21**(3), 427–444 (2009)
 38. Mišić, V., Aleman, D., Sharpe, M.: Neighborhood search approaches to non-coplanar beam orientation optimization for total marrow irradiation using IMRT. *Eur. J. Oper. Res.* **205**(3), 522–527 (2010)
 39. Lee, C.-H., Aleman, D., Sharpe, M.: A set cover approach to fast beam orientation optimization in intensity modulated radiation therapy for total marrow irradiation. *Phys. Med. Biol.* **56**(17), 5679–5695 (2011)
 40. Rocha, H., Dias, J., Ferreira, B., Lopes, M.: Selection of intensity modulated radiation therapy treatment beam directions using radial basis functions within a pattern search methods framework. *J. Global Optim.* **57**(4), 1065–1089 (2013)
 41. Djajaputra, D., Wu, Q., Wu, Y., Mohan, R.: Algorithm and performance of a clinical IMRT beam-angle optimization system. *Phys. Med. Biol.* **48**(19), 3191–3212 (2003)

42. Aleman, D., Kumar, A., Ahuja, R., Romeijn, H., Dempsey, J.: Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning. *J. Global Optim.* **42**(4), 587–607 (2008)
43. Bertsimas, D., Cacchiani, V., Craft, D., Nohadani, O.: A hybrid approach to beam angle optimization in intensity-modulated radiation therapy. *Comput. Oper. Res.* **40**(9), 2187–2197 (2013)
44. Dias, J., Rocha, H., Ferreira, B., Lopes, M.: Simulated annealing applied to IMRT beam angle optimization: a computational study. *Phys. Med.* **31**(7), 747–756 (2015)
45. Hou, Q., Wang, J., Chen, Y., Galvin, J.: Beam orientation optimization for IMRT by a hybrid method of the genetic algorithm and the simulated dynamics. *Med. Phys.* **30**(9), 2360–2367 (2003)
46. Li, Y., Yao, J., Yao, D.: Automatic beam angle selection in IMRT planning using genetic algorithm. *Phys. Med. Biol.* **49**(10), 1915–1932 (2004)
47. Lei, J., Li, Y.: An approaching genetic algorithm for automatic beam angle selection in IMRT planning. *Comput. Methods Programs Biomed.* **93**(3), 257–265 (2009)
48. Li, Y., Yao, D., Yao, J., Chen, W.: A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning. *Phys. Med. Biol.* **50**(15), 3491–3514 (2005)
49. Bangert, M., Ziegenhein, P., Oelfke, U.: Characterizing the combinatorial beam angle selection problem. *Phys. Med. Biol.* **57**(20), 6707–6723 (2012)
50. Bangert, M., Ziegenhein, P., Oelfke, U.: Comparison of beam angle selection strategies for intracranial IMRT. *Med. Phys.* **40**(1) (2013)
51. Ehr Gott, M., Johnston, R.: Optimization of beam directions in intensity modulated radiation therapy planning. *OR Spectr.* **25**(2), 251–264 (2003)
52. Lee, E., Fox, T., Crocker, I.: Integer programming applied to intensity-modulated radiation therapy treatment planning. *Ann. Oper. Res.* **119**, 165–181 (2003)
53. Yang, R., Dai, J., Yang, Y., Hu, Y.: Beam orientation optimization for intensity-modulated radiation therapy using mixed integer programming. *Phys. Med. Biol.* **51**(15), 3653–3666 (2006)
54. Schreibmann, E., Lahanas, M., Xing, L., Baltas, D.: Multiobjective evolutionary optimization of the number of beams, their orientations and weights for intensity-modulated radiation therapy. *Phys. Med. Biol.* **49**(5), 747–770 (2004)
55. Fiege, J., McCurdy, B., Potrebko, P., Champion, H., Cull, A.: PARETO: a novel evolutionary optimization approach to multiobjective IMRT planning. *Med. Phys.* **38**(9), 5217–5229 (2011)
56. Jia, X., Men, C., Lou, Y., Jiang, S.: Beam orientation optimization for intensity modulated radiation therapy using adaptive $l_{2,1}$ -minimization. *Phys. Med. Biol.* **56**(19), 6205–6222 (2011)
57. Robini, M., Zhu, Y., Liu, W., Magnin, I.: A stochastic framework for spot-scanning particle therapy. In: *Proc. Int. Conf. IEEE Eng. Med. Biol. Soc.*, pp. 2578–2581, Orlando, FL (2016)
58. Robini, M.: Theoretically Grounded Acceleration Techniques for Simulated Annealing, *Handbook of Optimization, from Classical to Modern Approach*, pp. 311–336. Springer (2012)
59. Robini, M., Reissman, P.-J.: From simulated annealing to stochastic continuation: a new trend in combinatorial optimization. *J. Global Optim.* **56**(1), 185–215 (2013)
60. Altobelli, E., et al.: Combinatorial optimization in radiotherapy treatment planning. *AIMS Med. Sci.* **5**(3), 204–223 (2018)
61. Hajek, B.: Cooling schedules for optimal annealing. *Math. Oper. Res.* **13**(2), 311–329 (1988)
62. Chiang, T.-S., Chow, Y.: On the convergence rate of annealing processes. *SIAM J. Control Optim.* **26**(6), 1455–1470 (1988)
63. Catoni, O.: Rough large deviation estimates for simulated annealing: application to exponential schedules. *Ann. Probab.* **20**(3), 1109–1146 (1992)
64. Robini, M., Magnin, I.: Optimization by stochastic continuation. *SIAM J. Imaging Sci.* **3**(4), 1096–1121 (2010)
65. Weber, U., Kraft, G.: Design and construction of a ripple filter for a smoothed depth dose distribution in conformal particle therapy. *Phys. Med. Biol.* **44**(11), 2765–2775 (1990)
66. Allison, J., et al.: Geant4 developments and applications. *IEEE Trans. Nucl. Sci.* **53**(1), 270–278 (2006)
67. Owen, H., Lomax, A., Jolly, S.: Current and future accelerator technologies for charged particle therapy. *Nucl. Instrum. Methods Phys. Res. A* **809**, 96–104 (2016)
68. Drzymala, R., Mohan, R., Brewster, L., Chu, J., Goitein, M., Harms, W., Urie, M.: Dose-volume histograms. *Int. J. Radiat. Oncol. Biol. Phys.* **21**(1), 71–78 (1991)
69. Scherrer, A., Küfer, K.-H., Bortfeld, T., Monz, M., Alonso, F.: IMRT planning on adaptive volume structures—a decisive reduction in computational complexity. *Phys. Med. Biol.* **50**(9), 2033–2053 (2005)
70. Martin, B., Bortfeld, T., Castañón, D.: Accelerating IMRT optimization by voxel sampling. *Phys. Med. Biol.* **52**(24), 7211–7228 (2007)
71. Catoni, O., Trounev, A.: Parallel annealing by multiple trials: a mathematical study, *Simulated annealing: parallelization techniques*, pp. 129–143. Wiley (1992)

72. Witte, E., Chamberlain, R., Franklin, M.: Parallel simulated annealing using speculative computation. *IEEE Trans. Parallel Distrib. Syst.* **2**(4), 483–494 (1991)
73. Sohn, A.: Generalized speculative computation of parallel simulated annealing. *Ann. Oper. Res.* **63**, 29–55 (1996)
74. Karger, C., Peschke, P.: RBE and related modeling in carbon-ion therapy. *Phys. Med. Biol.* **63**(1) (2018)
75. Thieke, C., Bortfeld, T., Küfer, K.-H.: Characterization of dose distributions through the max and mean dose concept. *Acta Oncol.* **41**(2), 158–161 (2002)

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